PLANET ALGORITHM

ARTIFICIAL INTELLIGENCE FOR A PREDICTIVE AND INCLUSIVE FORM OF INTEGRATION IN LATIN AMERICA

IDB
INTAL
Planeta

DIGITAL DIVIDENDS

I&T No. 44, Year 22, July 2018
ISSN 978-950-49-6426-1 | R.P.I. 5355458

Steering Committee / Antoni Estevadeordal and Gustavo Beliz.
Editorial Committee / Ana Inés Basco, Mariana Pernas, Cecilia Nicolini, Belisario de Azevedo, and Santiago Chelala.
Design / Andrea Pellegrino and Santiago Fraccarolli.
Proofreading (Spanish edition) / Yamila Rubbo.
Translations / María Inés Martiarena and Victoria Patience.

Beliz, Gustavo
Planet Algorithm / Gustavo Beliz. - 1st ed. -
Autonomous City of Buenos Aires: Planeta, 2018.
336 pp. ; 23x15 cm.

1. Technnology. I. Title.
CDD XXXXX

Impreso at Talleres Trama, Pasaje Garro 3160, CABA, in August 2018

Welcome

ANTONI ESTEVADEORDAL
Manager, Integration and Trade Sector,
Inter-American Development Bank

Latin America is facing new challenges when it comes to finding the best route to greater regional integration. Industry 4.0 can significantly help increase productivity and make our countries more competitive. At the same time, these changes pose questions about how technological progress will impact our productive matrix and export basket and how best to achieve consensus around regulations that will facilitate new forms of trade.

This issue of *Integration and Trade* focuses on a specific technology, one that is perhaps the most disruptive of the Fourth Industrial Revolution: artificial intelligence (AI). Machine learning processes are beginning to be used in food production, the automotive sector, finances, and services.

At present, nearly 90% of intraregional trade is tariff-free, but the regulatory architecture that has been built around the proliferation of preferential trade agreements (PTAs) is not giving companies the flexibility they need to compete in the digital economy. What form of regulatory framework for data exchange would be needed for integration 4.0 to enable us to move beyond the current patchwork of agreements and achieve a unified regional market? Can we expect a reshoring process whereby goods stop being produced in remote locations when labor costs become a less significant factor in determining the configuration of value chains? How can we prepare for this threat, given that intraregional trade has dropped by 26% over the last five years? Might AI, as a new factor of production, consolidate existing relations through greater efficiency, increasing flows of trade, products, knowledge, and information both globally and regionally?

We are living in a thrilling time of change. Image recognition has started to be implemented at border crossings to speed up customs procedures. AI is being used in ports to better organize container logistics. Other companies are using it for inventory management, to optimize the relationship between clients and suppliers, and in transportation using autonomous vehicles. Policymakers are even starting to access it during multilateral negotiations, which tend to be complex due to the enormous amounts of data they entail. These are just some examples of the many uses of AI that are explored throughout this issue.

AI is just another tool for reducing shortfalls in regulatory frameworks and improving the quality of institutions (in other words, the software of integration) and for helping to close gaps in infrastructure and physical connectivity (the hardware of integration) that have long characterized our region and made it less competitive than that of more developed countries.

The specialists that feature in this issue describe the current state of affairs and point to how AI could better support regional integration. At the IDB's Integration and Trade Sector, we are convinced that every step toward facilitating trade and reducing obstacles that impede trade in goods, services, and ideas will help improve Latin Americans' quality of life.

5-R/Es

THE FIVE R/EVOLUTIONS IN ARTIFICIAL INTELLIGENCE IN LATIN AMERICA

GUSTAVO BELIZ
Director, Institute for the Integration of Latin America and the Caribbean, Inter-American Development Bank

R/E. Revolution and evolution. In slow motion and at supersonic speed. Physical and digital. Visible and intangible. Among goods and among services. Built out of steel and built on data. On highways and in smartphones. In the 19th century and in the 22nd. Technophobia and techno-utopia… Latin America is beginning to feel the impact of technology-driven disruptions in which traditional models have not been entirely left by the wayside, but innovative new models have yet to fully take root. In this transitional scenario, the dramatic appearance of artificial intelligence (AI) as a new factor of production is pushing us to find new responses to the unfinished challenges of diversifying our economic patterns and adding value to our role in the global market in a way that is environmentally sustainable and socially inclusive.

Whether or not we know it, we now live on a new planet: Planet Algorithm. We are surrounded by algorithms, sets of logically organized operations that allow us to solve specific problems. With ever greater—exponential, even—frequency, cryptic formulas suggest what we should buy next, the film we should see this weekend, what our credit rating should be, and how likely we are to contract a certain illness. There is also a structural dimension to this individual impact, one that includes everything from the evolution of stock market transactions to commodity prices via mass transportation, legal decisions, and the industrial internet. We are witnessing the rise of imperceptible yet intelligent white-collar robots who are able to learn, unlearn, feel, and infer things, which makes them superior to the standard blue-collar robots that performed the routine tasks of the industrial era.

The scale and speed of data, together with ever vaster analytical capacities, is transporting us to a paradoxical reality in which our fears grow as technology becomes more efficient. Studies by INTAL/Latinobarómetro, Eurobarometer, and the Pew Research Center[1] show that over 70% of inhabitants fear that AI and robots will replace their current jobs and destroy more occupations than they can create. At the same time, apocalyptic visions of AI as a threat to humanity exist alongside developments that may lead to solutions to many of our major problems.

This publication takes a holistic, multidisciplinary approach to these questions, one that is grounded in solid empirical evidence. It brings together over 40 international experts to analyze the socio-economic implications of AI and put forward answers to the question of how our region can take best advantage of the opportunities AI offers in order to improve citizens' lives and integrate countries better while mitigating the risks that come with such massive change.

The research, reports, and case studies analyzed in the publication clearly show how AI is promoting a hybrid form of integration in which tariff negotiations still have a part to play but are no longer the sole focus of attention. Like geological layers, these now lie beneath more contemporary phenomena such as digital trade and e-commerce, innovation garages, and reshoring. The ways that we connect with the world and other countries are no longer solid or liquid—they are invisible, like algorithms.[2] It is precisely this intangibility, this absence of physical form, that we find unsettling: these new relationships take an entirely different form to the ones that we thought would exist forever, calling on us to redesign ourselves by analyzing the 5 R/Es, that is, the five revolutions and evolutions that I will outline here.

1 · THE R/EVOLUTION IN PRODUCTIVITY
From mass planning to the algorithm factory

International trade is being reshaped by two sweeping new trends.[3] In the financial sphere, the hardening of liquidity conditions for emerging countries and the stress caused by the increased cost of money in the world. In the real economy, new technologies and Industry 4.0, which now demand permanent innovation to prevent the goods and services that are exported from becoming obsolete. These two trends merge into a single concept: selectivity. Just as capital flows and investments will become more selective, consumers and companies will need to compete in increasingly demanding global markets.

Given this context, business models in most sectors are being rapidly transformed. AI affects both production techniques within companies and relationships with customers and suppliers. New goods, industries, and services are emerging which are forcing us to redefine our export matrix and set aside the old parameters of concentration. Stanford's (2017) speedometer for AI shows that the rate of these changes is dizzying and exponential. Today, 54% of high-level executives at global companies are investing in AI[4] and are thus effectively building the global algorithm factory. This includes every sector of the economy and is the driving force behind real-time cyberphysical production.

According to Purdy and Daugherty (2016), AI can increase labor productivity by up to 37% when human labor is boosted by machine learning tools. For Latin America, the economic benefit of this massive increase in productivity has been estimated at US$700 billion.[5] This is reflected in greater efficiency, speed, new personalized and prototyped products,[6] and new consumer, investment, and trade platforms[7] that are based on the concept of predictive integration, wherein the demands and needs of other players can be more clearly envisaged, regardless of whether they are national or local institutions or individual consumers.

Indeed, as several authors observe in this publication, AI is beginning to be understood as a new factor of production in itself (an amalgam of physical capital and labor), one that might outperform human intellectual capacities, learning more and learning faster, and giving rise to a new virtual player when it merges with human skill in what is referred to as "cobotization."

Latin America must not be left behind in this race. Any effort to reduce barriers to trade and facilitate the internationalization of our companies—particularly the SMEs that make up the base of the business pyramid—would be a step in the right direction. During this process, AI should be seen not as an option but instead as an obligation. The most immediate threat is not applying AI but rather overlooking it.

2 · THE R/EVOLUTION IN INFRASTRUCTURE

From highways to the cloud

Next, we need to close the infrastructure gap and take an exponential trade leap based on an up-to-date agenda. While we are working on bringing down logistics costs through all the old channels, developed countries are gaining ground on us. In Normandy, for example, in the north of France, and along stretches of Route 66 in the United States, solar highways are enabling energy to be gathered that can then be used to power city streetlights, recharge electric car batteries, and change the direction of traffic depending on demand, using AI sensors.[8] This speeds up freight transportation and reduces losses of goods caused by unnecessary delays.[9]

By cutting down on red tape and transforming ports into smart corridors that are better placed to become part of logistics chains, AI is improving the competitiveness and operational efficiency of the sectors that were ahead of the game in implementing it.

Infrastructure is the scaffolding used by enterprises from all manner of economic sectors to scale up their operations. Without digital infrastructure, agricultural areas that are far from major cities will be unable to apply precision agriculture techniques that use AI to find the best combination of inputs for each geographic area or to simulate growing scenarios, detect pests using image recognition, or focus pesticide use exclusively on affected plants.[10]

In the manufacturing industry, vehicle companies are using lidar (light detection and ranging) technology in combination with AI and big data to make vehicles fully autonomous. Some surprising advances in this field have taken place in Latin America as a result of the partnership between Embraer and Uber to manufacture flying cars. But are we getting highway infrastructure ready for new forms of transportation like these?

AI is also being used to create smart grids that coordinate energy generation with real-time demand, which saves money and reduces environmental impact.[11] IT services based on AI, big data, and cloud computing are another key area for export diversification, driven by knowledge-based services that need state-of-the-art infrastructure.[12]

Drones, which are already being used in pilot tests for B2C (business-to-consumer) deliveries, could make up an intelligent aerial network that might not entirely replace today's land transportation system but may nonetheless constitute a new, faster, more efficient form of shipping, particularly in congested areas, remote locations, or hostile environments. Automated deliveries of this type would cut down on operating costs by 60% due to route optimization.

Technology can to reset the development counter, even when it comes to physical connectivity. Infrastructure 4.0 is a fast-track to closing the gap between Latin America and more developed countries.

3 · THE R/EVOLUTION IN GOVERNANCE

From machine learning to government learning

Automation is not about putting governments on autopilot. Far from it. Many of Latin America's trade partners or competitors in Asia and Europe have launched national strategies to introduce AI-based tools into their economies and increase the productivity of tradable sectors. This is true of China, whose AI development plan sets out to increase energy efficiency by 10%. It is also true of India, which is planning to be the global market leader in 5G technology, software, and ICTs. Another example is Japan, which is seeking to find solutions its aging population using AI applications in the care economy and the health services industry.[13]

The governance of AI needs to be based on solid incentives, rewards, and punishments. Examples of rewards include initiatives which aim to increase research into AI, as is happening locally in Canada through the Pan-Canadian Artificial Intelligence Strategy, through which specialized centers receive public funds to fast-track processes and applications. On the punishment side, AI governance needs to come down hard on those who break laws on personal information and privacy, seek to make money through cyberpiracy, or implement biased algorithms that reproduce prejudice or affect people's lives or possibilities of progress.

Spain has also launched a public consultation process to shape its digital strategy, which rests on five pillars: data ethics, the creation of digital ecosystems, regulation, technological infrastructure, and digital employment. The tourism sector, the keystone of the Spanish export sector, has been identified as one of the key application areas. An expert council has also been created to draft a white paper on the guiding principles for AI use in the country. The council is made up of representatives from academia, the private sector, and civil society and depends on the office of the Secretary of State for Information Society and the Digital Agenda.

The British government has pledged funds for the creation of a Centre for Data Ethics and Innovation, a public organization whose role will be to provide public policy and regulatory recommendations and to carry out continuous monitoring to guarantee security around AI use. Other countries have their own Industry 4.0 programs that bring AI together with local production.[14] This is a debate that Latin Americans have yet to fully engage in.

Why aren't we creating a network of regional experts on AI? This could function as a space in which to decide which public policies are most in line with the realities of life in our region, and through which to articulate different efforts and initiatives. A lot can be done from the public sector to lighten the burden of labor transition costs during the automation process, promote the teaching of soft skills within the education system, and monitor global trends to anticipate changes in demand for labor. Given the significance of Latin American states' current public investments in human capital, AI could provide tools for modernizing defense and security forces, supporting our teachers, and facilitating the work of medical personnel in clinics and hospitals. Other good examples of areas with potential for AI applications are the Mercosur Digital Initiatives and focused government procurement to promote innovation.

We need to come up with a strategy of our own. We need to ask ourselves how we are going to use AI to add value to our strategic sectors in partnership with states that support and foster the private sector. It is time to move from machine learning to government learning and to leave behind the notion of the public sector as languishing in routine-based silos. Instead, we have to embrace the idea of a creative, dynamic, neural network–based government.[15]The boxes on the geopolitical algorithm board will fill up sooner rather than later. We already know all about the price that countries pay for lingering on the platform while they watch the development train steam off into the future. Inaction is a decision like any other: far from neutral, it brings benefits and losses.

4 · THE R/EVOLUTION IN EQUALITY

From augmented reality to augmented humanity

We are witnessing a three-dimensional robotization process that is simultaneously creating, replacing, and destroying jobs. We mustn't deceive ourselves but should instead be preparing for the changes ahead through more and better education, new regulatory frameworks, better-quality institutions, and more inclusive agreements. Just as some jobs will inevitably disappear, be it as a consequence of innovation or through mere shifts in preferences, new professions will emerge in which humans and machines work together in an unprecedented way. Likewise, different fields are coming together to form new educational paths: in the United States, for example, the University of Illinois now requires that applicants to its postgraduate medicine program first have an engineering qualification.

AI may also be a source for the creation of new, as yet nonexistent jobs, such as experts in vertical agriculture, cobotization trainers, FinTech planners, cyberspace anthropologists, auditors for the sharing economy, virtual reality travel planners, and data training consultants, among many others.[16]

Similarly, today's professions may start to improve through interactions with increasingly sophisticated and personalized AI programs, as is described in many of the examples included in our report: from public policy consultants to judges and lawyers, from agronomists to owners of MSMEs, from urban designers to creative industry workers, from multilateral agreement negotiators to experts in cybersecurity and violence prevention.

Many pundits have heralded the end of work. However, Austria, the United States, and Germany, countries with high numbers of robots per inhabitant, have low unemployment rates or lost fewer jobs than other countries during the last financial crisis. Of the 535 most common jobs in the world, automation only starts to feature as a threat to number 21 on the list (Tegmark, 2017).

We are entering an age in which regular employment and salaries are a thing of the past, in which the automation and digitalization of the economy are encouraging indirect labor relationships, such as freelancing and telecommuting, which have very different implications to those of traditional jobs. In the innovation ecosystem, this new state of affairs is common currency. Start-ups need to be flexible and take a trial-and-error approach, which the rigidity of classic labor markets makes difficult.

But just as there are new forms of work, there are also new forms of slavery. The absence of institutions and rules to legitimize the relationship between employers and employees make the situation ripe for different types of exploitation and abuse entailing different degrees of violence, which inevitably affect the most vulnerable parts of society.[17] At the same time, as several articles in this publication report, AI is opening up vast possibilities in the fields of education, health, and legal services, broadening the potential horizons for social inclusion. These include pilot programs for individualized, algorithm-based learning experiences that show how AI can help boost school exam scores by 15%. The cost of sequencing genomes has fallen five times faster than predicted by Moore's Law, making precision medicine more accessible. Watson's services help detect illness early through automated image recognition. We can even cut down on the time we waste on paperwork by up to 75% by applying AI to government processes.

Perhaps the most critical question of our time is whether the mass spread of AI will lead to a utopian state of affairs in which machines take care of the most boring and dangerous tasks while human beings focus on doing more creative work and enriching community life. Or, in contrast, will power be more concentrated in the future and inequality more profound, as part of a new economy that revolves around a handful of superstars?[18] We can't just sit back and watch. We can't even allow ourselves the benefit of the doubt. We need to create the necessary conditions for guaranteeing that the innovation process is more likely to lead to the former scenario than the latter, so that we can leave the risk of technological discrimination behind. The point is using digital dividends to promote equality in the most unequal continent on earth.

5 · THE R/EVOLUTION IN ETHICS

From big data to big values

It may sound paradoxical, but if Latin America wants to put humans at the center of its concerns, it has no choice but to design public policies around robots. The region needs to heed the calls of the inventors, investors, and visionaries who want to move from implicit philosophical discussions to explicit legal and regulatory decision-making that anticipates the ethical consequences of AI sweeping through our daily and productive lives. The rest of the world has already started this process.

AI can imitate human creativity and make music, poetry, or visual art. It can identify emotions based on voice recognition, make decisions when it is asked to analyze precedents for a legal decision, or outperform the most precise of surgeons when used in an operating theater. Achievements that were until recently the stuff of sci-fi films are now becoming a reality.

However, just as advances in AI surprise us, they also beg new questions. Cathy O'Neil, author of the best-selling *Weapons of Math Destruction*, warns that although algorithms may seem scientific and objective, they are actually highly subjective and are little more than "opinions embedded in code." In other words, they are no better than we are. They are modeled on us.

Algorithms use inductive logic. They analyze past information and make predictions based on it, which thus run the risk of maintaining the status quo. Are the enormous inequalities, prejudices, and gender inequalities in Latin America something that we want to perpetuate? The answer, clearly, is no. We need to change things. Algorithms put us to the test and challenge us to improve our values and try to avoid bias when we look at the formulas behind them.

Their basic input, data, also needs to be carefully checked. Academia and civil society are one step ahead of public policymakers in this debate. There are five manifestos that have been signed by over 12,000 scientists and global experts that warn of the risks AI poses (to privacy, national security, and transparency) and set out ethical standards to ensure that AI does not accentuate inequality.

The Institute of Electrical and Electronics Engineers (IEEE) has published a manifesto containing ethical principles to prevent autonomous systems from violating basic human rights and damaging the environment. The Copenhagen Letter, signed by many high-profile scientists, invites readers to share in a vision of progress that moves beyond the concept of innovation by reminding them that human beings must be at the heart of any technological action: "tech is not above us," the text stresses. The ethical manifestos published by the Future of Humanity Institute, Oxford University, and the Future Society, which is backed by Harvard, point out the risks of military applications of AI and advocate for a form of global governance that will provide an ethical framework for this technology. The Open Letter published by the Future of Life Institute was signed, among others, by Stephen Hawking, and seeks to promote interdisciplinary AI studies because its spread will have enormous legal, economic, social, and philosophical impacts. The perspective needs to be a holistic one.[19] Nor is the private sector just sitting back and watching. Tech giants IBM, Amazon, Facebook, Google, DeepMind, and Microsoft have created the Partnership on AI to Benefit People and Society. Its objectives are to empower as many people as possible around the use of AI tools and to be part of the debate about its legal, ethical, social, and economic consequences. There are profound dilemmas around human rights, privacy, and representation.

The aim of Isaac Asimov's three laws of robotics was to prevent humans from being physically harmed by robots. It is up to us to create new rules for cobotization, broader legal and ethical cannons that ensure that human dignity is not affected by the loss of jobs to combinations of steel and algorithms.

If we lay down solid foundations, AI's potential is thrilling. For example, Broussard (2018) explains how it could be used to solve problems that are endemic to our societies, such as the lack of transparency in political campaign finance.[20] AI may also prove useful in evidence-based policy-making, notably at the stage of identifying the most appropriate form of intervention, implementing a program, and measuring the desired impacts.[21]

Latin America needs to start relying on the human factor, on talent. Passion, commitment, sacrifice, teamwork, and creativity. The areas where we outperform machines. We cannot measure our successes and failures only in terms of productivity nor can we marvel at ingenious gadgets if the digital dividends from them are not channeled into making us better individuals and if technology does not translate into public policies that improve the vast majority of people's lives.

Rising to this challenge is, of course, possible, if our starting point is a technological humanism that puts people at the core of its efforts. On Planet Algorithm, we need to do more than just improve our abilities to predict things—instead, we need to make everything much more predictable and to create a digital New Deal that guarantees social inclusion. Intelligent integration is about so much more than just algorithms.

NOTES

[1] A study by the Pew Research Center found that 76% of people in the US believe that financial inequality will increase with automation. For more on this, see Smith and Anderson (2017) and Basco (2017).

[2] In a thought-provoking book, Rubin (2018) examines these changes in the way people connect to one another and the adulteration of concepts of intimacy and privacy. The author lays out different scenarios of how our relationships with the world could be modified through developments in virtual reality technology.

[3] Mesquita Moreira (2018) analyzes the uncertainty that currently reigns in the world of foreign trade. The publication explores the uncertainty that currently characterizes foreign trade and recommends specific actions to make regional integration a reality and a vehicle for citizen well-being.

[4] See the article in this issue by Rao.

[5] See the article in this issue by Ovanessoff and Plastino for a regional approach to the relationship between AI and productivity, and Brynjolfsson, Rock, and Syverson (2017) for a global perspective.

[6] The challenges entail capturing the value generated by product use and moving from traditional products to product platforms, manufacturing intelligent products that include services, and prioritizing access to products over ownership. Through open innovation platforms, cooperation mechanisms are being established between companies that allow

them to accelerate the results of R&D&I initiatives (see INTAL, 2018a).

[7] Salesforce (2017) shows the changes that AI can bring to the world of business and the growing receptiveness to it among firms and consumers.

[8] Gesing, Peterson, and Michelsen (2018) describe AI's latest contributions to logistics.

[9] Each additional day of delay reduces trade by 1% and by as much as 7% in the case of perishable products.

[10] García Zaballos and Iglesias Rodríguez (2017) highlight the role of new technologies in exporter SMEs from Latin America.

[11] A system of this type has been implemented by Agder, the third largest energy company in Norway.

[12] INTAL (2018b) looks at KBSs in the region.

[13] Acemoglu and Restrepo (2017) analyze the potential of new technologies for solving the problems that come with aging populations.

[14] Portugal, Germany, United Kingdom, Spain, The Netherlands (Smart Industry) and France (Alliance Industrie du Futur), and many others. In 2018, the White House created two special committees on AI in the US, one on machine learning and another on R&D in AI.

[15] Domingos (2018) includes a map of the impact of machine learning on different productive sectors.

[16] Cognizant (2017) describes 21 new professions that will play a vital role in the economy from 2028 onward.

[17] Graeber (2018) examines new trends in the world of work and the problems these are causing to community life.

[18] The dilemmas around income distribution are analyzed in this publication by Anton Korinek.

[19] Different aspects of the impact of AI are analyzed by multidisciplinary experts in Brockman (2015). The birth of the pro-ethics movement in this field is described in detail by one of its protagonists in Life 3.0 (Tegmark, 2017).

[20] In 2016, the author launched www.campaign-finance.org, which uses AI to facilitate the cross-referencing of information on political campaign financing from the private sector and different parts of the state.

[21] For more on AI and its impact on evidence-based policy-making, see Bhatt et al. (2016). The authors describe methods wherein AI allows contextual variables to be better controlled, ranging from individual perception to the presence of other individuals, and so on.

REFERENCES

Acemoglu, D., and Restrepo, P. 2017. "Secular Stagnation? The Effect of Aging on Economic Growth in the Age of Automation." *American Economic Review* 107 (5): 174–179.

Basco, A.I. 2017. "Techno-Integration in Latin America. Institutions, Exponential Trade, and Equality in the Age of Algorithms." Technical Note No. IDB-TN-1340. Inter-American Development Bank.

Bhatt, M., Suchan, J., Schultz, C., et al. 2016. "Artificial Intelligence for Predictive and Evidence-Based Architecture Design." In: Proceedings of the 30th AAAI Conference on Artificial Intelligence. Association for the Advancement of Artificial Intelligence.

Brockman, J. 2015. *What to Think About Machines That Think: Today's Leading Thinkers on the Age of Machine Intelligence.* New York: Harper.

Broussard, M. 2018. *Artificial Unintelligence: How Computers Misunderstand the World.* Cambridge: The MIT Press.

Brynjolfsson, E. Rock, D., and Syverson, C. 2017. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." NBER Working Paper No. 24001. Cambridge: NBER.

Cognizant. 2017. *21 Jobs of the Future.* Center for the Future of Work.

Domingos, P. 2018. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World.* New York: Basic Books.

García Zaballos, A., and Iglesias Rodríguez, E. 2017. *Economía digital en América Latina y el Caribe. Situación actual y recomendaciones.* Washington DC: Inter-American Development Bank.

Gesing, B., Peterson, S.J., and Michelsen, D. 2018. *Artificial Intelligence in Logistics.* Troisdorf: DHL Trend Research.

Graeber, D. 2018. *Bullshit Jobs: A Theory.* New York: Simon & Schuster.

IFR. 2018. *Annual Report.* International Federation of Robotics.

INTAL. 2018a. "Industria 4.0." IADB Technical Note. Mimeograph.

INTAL. 2018b. "Los servicios basados en conocimiento: ¿el futuro del desarrollo?" *INTAL Connection* No. 260, April, INTAL, Inter-American Development Bank.

Mesquita Moreira, M., editor. 2018. "Connecting the Dots. A Road Map for a Better Integration of Latin America and the Caribbean." Special Report. Integration and Trade Sector, IDB.

O'Neil, C. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy.* New York: Crown Ed.

Purdy, M., and Daugherty, P. 2016. "Why Artificial Intelligence Is the Future of Growth." Working Paper, Accenture.

Rubin, P. 2018. *Future Presence: How Virtual Reality Is Changing Human Connection, Intimacy, and the Limits of Ordinary Life.* New York: HarperCollins.

Salesforce. 2017. *The AI Revolution: Insights into the Next Era of Customer Relationships.* Special Report. Salesforce Research.

Smith, A., and Anderson, M. 2017. "Automation in Everyday Life." Pew Research Center. Washington, DC. October 4.

Stanford. 2017. *The IA Index 2017 Annual Report.* Stanford University. Available at https://aiindex.org/

Tegmark, M. 2017. *Life 3.0: Being Human in the Age of Artificial Intelligence.* New York: Alfred A. Knopf.

PLANET ALGORITHM: 10 KEY POINTS

· 1 ·

AI is a new factor of production which Latin America needs to invest in to improve economic growth in the region and foster integration 4.0.

· 2 ·

Putting a techno spin on trade diplomacy, using deep learning and next-generation big data, could breathe new life into current trade negotiations and the region's strategic role in global value chains.

· 3 ·

AI will allow us to build more sophisticated predictive regional integration scenarios using an innovative set of advanced analytical tools.

· 4 ·

AI-based cognitive infostructure—or an intangible, infrastructure 4.0—is an opportunity to close the physical connectivity gap in Latin America.

· 5 ·

AI has burst onto the scene in the region at a time when the risk of job automation stands at 39%, posing challenges and opportunities that need to be examined from a humanist perspective.

· 6 ·

Cobotization, the convergence of AI and new "digital workers," has sparked a growing, urgent demand for retraining and reskilling.

· 7 ·

The application of AI to global services poses the question of how we can add value and diversify the economy of a region that specializes in tasks that could potentially be codified.

· 8 ·

A form of algorithm-based collective intelligence could prevent an AI "rebellion" by anticipating the ethical risks of handling, analyzing, and producing large quantities of data.

· 9 ·

A neural network–type government that uses AI to improve social wellbeing would promote a fair distribution of digital dividends.

· 10 ·

It will be essential to build a Latin American framework for AI based on strategic priorities that capture innovative advantages and allow us to diversify the region's productive matrix.

1 AI is a new factor of production which Latin America needs to invest in to improve its economic growth and foster integration 4.0.

AI is more than just a new technological trend. It is a unique hybrid of capital and labor that creates a completely new productive workforce that can learn by itself. In the medium term, the growth rate for the region's GDP could go from 3% to 4% based solely on this factor. Almost 50% of that increase could be generated by an increase in productivity, as AI enables human workers to focus their efforts on the work that adds the most value. However, it is estimated that the impact of AI on the GDP of Latin America and the Caribbean (LAC) will be up to three or four times greater than in developed economies. This will make closing the development gap even harder than is already the case and it reveals the risks that come with delaying the process of creating the conditions necessary for consolidating AI and using it to increase productivity. This exponential change is unfolding simultaneously in different sectors that are key to Latin America's productive and export profile. For example, image recognition is used in precision agriculture to target fumigation, increasing yields per hectare by up to 30% through AI. AI is also being applied in the healthcare sector to provide diagnoses that are up to 96% accurate. In the automotive industry, it is estimated that 37% of trips will be made by autonomous vehicles by 2030. In e-commerce, it takes the form of chatbots that can detect emotion and provide 24/7 customer service at no extra cost. New global services are emerging and value chains are being reshaped based a mix of material goods and data.

ESTIMATED VARIATION IN GDP IN 2035

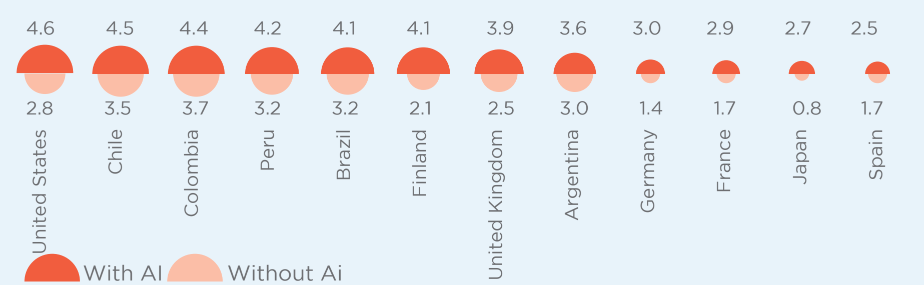

2 Putting a techno spin on trade diplomacy, using deep learning and next-generation big data, could breathe new life into current trade negotiations and Latin America's strategic role in global value chains.

AI is intensifying trends that have already emerged in international trade patterns, prompting changes in global value chains and trade and investment flows and leading to the reshoring of tasks that can be codified.

It could also prove a powerful way to accelerate the process of negotiating and managing trade policies.

The analysis of huge flows of data on trade volumes, tariff positions, technical and sanitary standards, impact assessments, rules of origin, and feasibility estimates could affect both the duration of multilateral talks and the political factors underlying them. The use of AI applications could promote negotiations that require consensus between multiple countries by making the negotiation process more efficient: AI increases the ability to process huge amounts of data on existing trade agreements and can suggest new paths that are in line with ongoing negotiations.

Automated data analysis allows us to track several years of decisions, which can help us understand the reasons underlying other parties' arguments and assess whether or not they are still valid. When objective matters can be solved using documented information that is tracked intelligently, negotiations speed up considerably. Prototypes are already being applied to ongoing negotiations, such as the Mercosur–Canada talks: this is happening from Brazil through the Intelligent Tech & Trade Initiative (ITTI), which has received support from private tech firms and is to be presented at the World Trade Organization (WTO).

3 AI will allow us to build more sophisticated predictive regional integration scenarios using an innovative set of anticipatory analytical tools.

A comparative efficiency exercise has shown that AI models can improve predictive capacity by up to 300%. Predictions based on artificial neural networks (ANNs) have proven highly useful for establishing connections between variables that may change suddenly. Using this method to predict changes in the agricultural commodities market yields results that are far more accurate than traditional econometric models. In financial and capital markets, algorithms predict investors' risk profiles with 95% accuracy rates.

The geographic proximity between different variables can be a decisive factor in trade, so image recognition and patter identification technology have astonishing predictive potential. Macroeconomic policy design will be increasingly influenced by invisible assumptions and undefined expectations and it will depend more and more on people's real activity through the traces they leave in their online activity. This predictive power increases yet more when it merges with other scientific disciplines, such as cloud computing or neuroscience.

4 AI-based cognitive infostructure—or an intangible, infrastructure 4.0—is an opportunity for closing the physical connectivity gap in Latin America.

Classic infrastructure for ports, highways, dams, roads, and logistics will be essentially transformed by this digital explosion. AI has concrete applications in the world of transportation, where it can be used to assign containers port space in real time and to optimize inventory management through sensors and the Internet of Things.

Cloud computing enables access to a global data deluge, 80% of which is unstructured. AI can reduce energy consumption by up to 10% through intelligent grids that automatically adapt to supply and demand. It can also be used to read traffic signs and urban markers to optimize mobility within an interconnected transportation network that includes electric, autonomous, and shared cars. Open-source maps and sensors

that predict the behavior of crowds of pedestrians can be used to adjust a robot's path in real time, enabling a kind of "socially aware navigation."
Cognitive infostructure strengthens physical infrastructure to create intelligent logistics corridors based on agile data flows. Real-time information combined with the digital universe are the building blocks for the much-needed reinvention of trade facilitation.

STAGES OF EVOLUTION IN THE MARITIME-PORT INDUSTRY

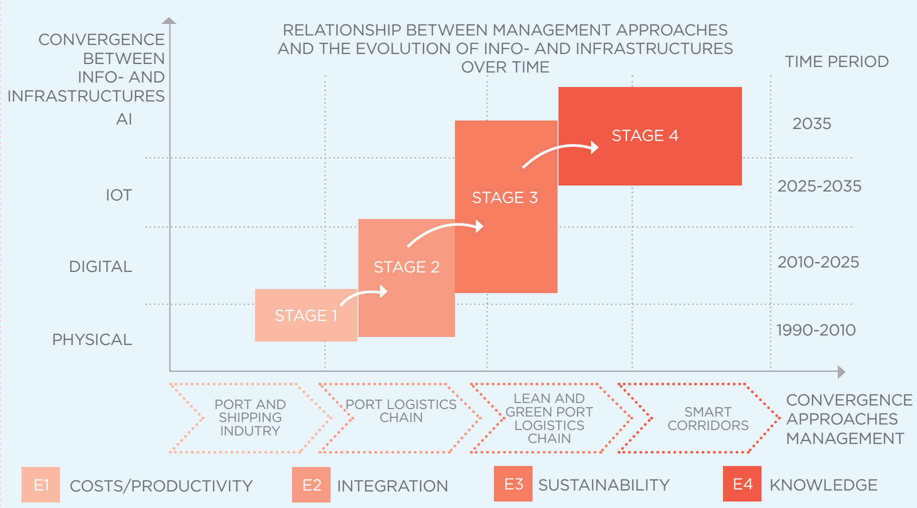

5 AI has burst onto the scene in the region at a time when the risk of job automation stands at 39%, posing challenges and opportunities that need to be examined from a humanist perspective.

Over 70% of Latin Americans, Europeans, and North Americans think that AI and robotics are a threat to employment. According to estimations, the risk of job loss in a given country as a consequence of automation fluctuates between 65% and 10%, yet there are also predictions that suggest that only a few tasks within each job can be automated entirely.
Planet Algorithm presents a new way of measuring this risk, one that takes aggregate data into account and thus allows the risk of automation to be tracked over time. The Synthetic Indicator for Automation Risk contemplates different aspects of this phenomenon: education levels, productive structure, robot numbers per industrial worker, the extent of ICT use, and the software content of exports. By including different socio-economic factors as part of a dynamic monitoring process, the average risk of job automation in Latin America is 39 on a scale of 1 to 100, where 100 is the maximum risk. The figures for each country in the region range between 36 and 43. The lower the country's per-capita GDP and the higher its income inequality, the greater the risk of job automation.
Designing public policies to carefully manage the technological transition of displaced workers into new jobs will thus become essential. Private-sector involvement in this process is a priority: today, 54% of high-level executives at global companies are investing in AI and automation, which will affect every sector of the economy.

RISK OF AUTOMATION FOR 37 SELECTED COUNTRIES (IN %)

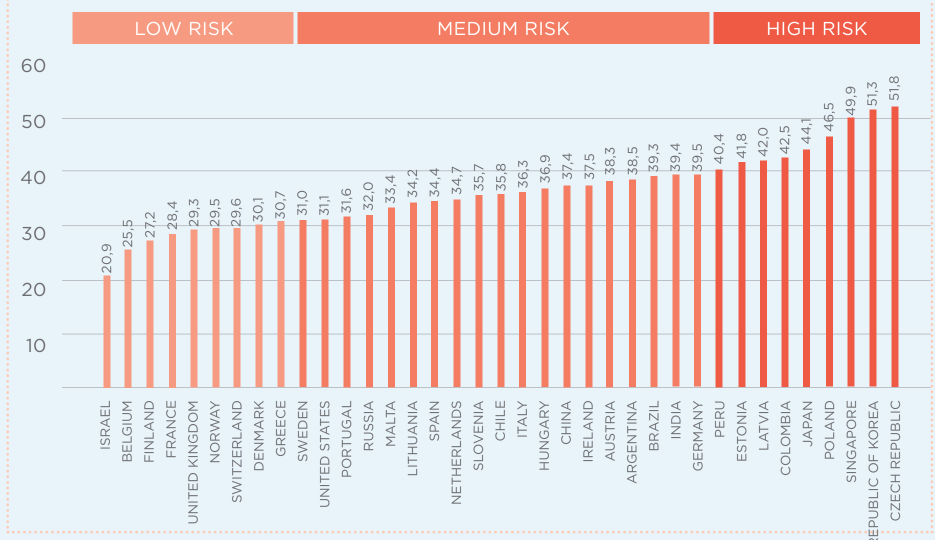

6 Cobotization, the convergence of AI and new "digital workers," has sparked a growing, urgent demand for retraining and reskilling.

Cobotization can prevent job loss by promoting the creation of a workforce that uses augmented intelligence, wherein AI raises the limits to traditional capacities. With the rise of so-called digital workers, nearly three-quarters of the impact of automation unemployment will take place within current job descriptions. By 2030, it is expected that workers will spend two hours less per week on routine, automatable tasks, which will allow them to concentrate on more complex, interactive ones.
They will spend twice as much time as they currently do on problem-solving and 41% more on critical thinking and reasoning. They will use verbal communication and interpersonal skills 17% more often per week and need to develop a stronger entrepreneurial mindset.
Although it is less fatalistic than the job replacement scenario, the idea of human workers and AI operating alongside one another in the region is not without challenges. Many people will need to reskill, absorb new knowledge halfway through their careers, and pay heed to the growing importance of in-

novation and data and digital skills. At the same time, new professions will emerge from this new productive conversation with AI, as is already the case with experts in clean energy and the green economy, an area where jobs are being created three times faster than in traditional fields.
Other examples include those who design and construct smart ecobuildings using new materials, new techniques for manufacturing autoparts for the transportation industry which use both data and steel, designers of synthetic vertical architecture and the bioeconomy, or social workers within the care economy. To favor the development of the labor market, there need to be incentives to boost public and private investments that strengthen basic research capacities around AI and automation, including robotics, autonomous systems, deep learning, and quantum computing.

SKILLS REQUIRED BY EMPLOYERS (% VAR. 2012–2015)

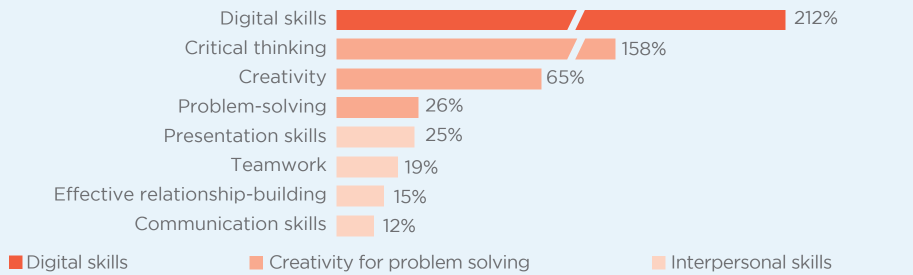

7 The application of AI to global services poses the question of how we can add value and diversify the economy of a region that specializes in tasks that could potentially be codified.

E-commerce and the offshoring of knowledge-based processes are global markets worth US$183 billion and they account for 20% of LAC exports. The challenge for economies that compete in the global services market is moving from codable tasks (like accounting, legal consultancy, and call centers) to tasks requiring high levels of creative intelligence, such as R&D and the development of software and new technologies, while also taking advantage of the servitization opportunities that are associated with local productive structures.
In Chile, for instance, a firm has developed an algorithm that analyzes animal protein-based foods and generates vegan alternatives with greater nutritional value. A Mexican start-up offers AI services for optimizing large companies' inventories. Developments in Argentina include drones that can react to unforeseen situations and 3D models of grapevines that allow crops to be monitored in real time and corrective actions to be undertaken. These bots are used by numerous companies in the region for online trade operations, targeted marketing for online sales, and support services for customers and potential buyers as part of sales and postsales process.

8 A form of algorithm-based collective intelligence could prevent an AI "rebellion" by anticipating the ethical risks of handling, analyzing, and producing large quantities of data.

So-called black box algorithms could give rise to biased decisions or facilitate access to higher-quality goods and services. To prevent the former and encourage the latter, we need spaces for multilateral dialogue that guarantee an inclusive form of governance for AI. Over 12,000 scientists and entrepreneurs have already warned of the risks of leaving decisions in the hands of autonomous processes that are not appropriately supervised. The world's main research centers see this as a top priority, as is evidenced by the creation of the Algorithmic Justice League at MIT, the rise of the Institute of Electrical and Electronics Engineers as the world's leading engineering center, and the ethical manifestos and core values expressed by the Copenhagen Letter, the Future of Humanity Institute at Oxford University, the Future Society think tank at Harvard, and the AI Open Letter published by the Future of Life Institute.

Partly inspired by governance initiatives like the United Nations Intergovernmental Panel on Climate Change, these publications seek to define criteria for binding and nonbinding legislation (hard and soft law) around AI use. The articulation of consensus is the most effective weapon for combating the new risks that come with AI, such as cybersecurity threats from hackers or the risk of inappropriate algorithm use, such as might happen with driverless vehicle technology being used as an autonomous weapon. Other additional risks include control issues and social, economic, and ethical matters that urgently need to be addressed. The tide of data that makes up our digital fingerprint needs to be organized based on principles of quality, reliability, access, and transparency.

RISKS ASSOCIATED WITH ARTIFICIAL INTELLIGENCE

1. PERFORMANCE RISKS	2. SECURITY RISKS	3. CONTROL RELATED RISKS	4. ETHICAL RISKS	5. ECONOMIC RISKS	6. SOCIAL RISKS
Error	Hacking	AI "rebellion"	"Lack of values"	Job displacement	Spread of autonomous weapons
Bias	Privacy	Inability to control malicious AI	Harmonizing values	Concentration of power in a winner-takes-all scenario	"Intelligence gap"
Lack of transparency or the "black box" effect	Open-source software		Harmonizing goals	Responsibility	
Explainability				Reputation	
Performance					

9 A neural network–type government that uses AI to improve social wellbeing would promote a fair distribution of digital dividends.

AI encourages winner-takes-all scenarios, in which a small number of superstars reap massive benefits. This exacerbates the inequalities within economies and widens the gap between countries: if just a handful of nations are able to develop superstar firms or initiatives, the remainder will suffer constant reductions in their terms of trade. Algorithms can lead to biased decisions and replicate erroneous human behavior. A form of AI that promotes social wellbeing will need to be built on public-private partnerships and knowledge transfers between academia, the state, and the business world. To achieve this, the region needs to work in a spirit of experimentation and prototyping for state policies that allow the creation of an intelligent public sector that will generate greater social impact.
AI tools that have been developed in the region will bring down the cost of remedial school programs by 40% and increase access to specialized content by 25%, providing teachers with support and adapting contents to each student. Thanks to innovation through deep learning, healthcare professionals can spend less time on data collection and routine tasks and more time on patients themselves. In the legal system, AI is being used to predict outcomes, draft verdicts, and settle cases in just 20 seconds, with an average accuracy rate of 96%. AI used in combination with brain–computer interfaces promises exponential changes in the way we communicate with objects. Nobel laureates have stressed the role of AI for processing information in real time and making government actions more predictable, including for optimizing evidence-based policy-making to achieve more precise impact analyses. With this aim in mind, the first step is generating data standards, as is already happening with electronic clinical histories or personalized education initiatives, which are increasing exam pass rates by 15%.

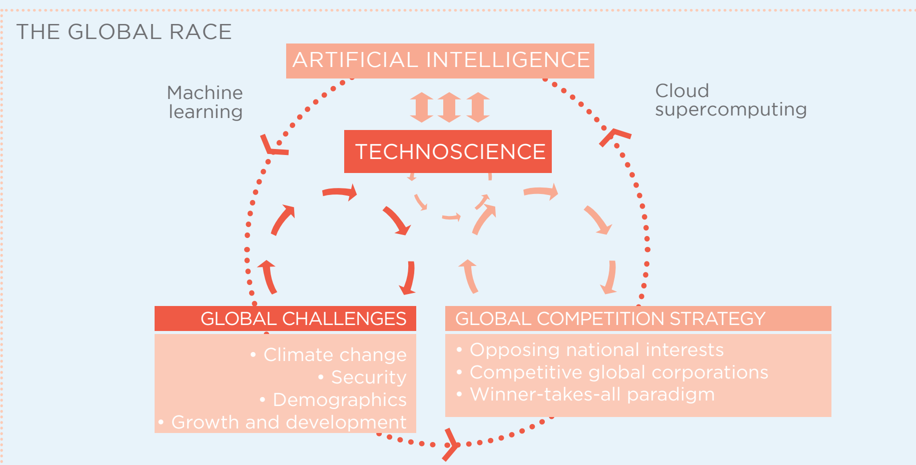

10 It will be essential to build a Latin American framework for AI based on strategic priorities that capture innovative advantages and allow us to diversify the region's productive matrix.

Over 13 developed and developing countries have made progress on plans for AI to strengthen the sectors that they believe to be key to their economies and the social wellbeing of their populations. This is the route India has taken in connection with 5G ICTs, China in connection with energy, and Japan in connection with healthcare. Spain has brought together a council of experts from academia, the public sector, and the private sector to draft a white paper on AI and public policy, with a focus on tourism, which is a key sector in its balance of trade. In the US, the White House has created two special committees on AI, one on machine learning and another on R&D in AI.
These strategies set out ambitious plans for investment in R&D, incentivizing training schemes to create experts, generating value, increasing export capacity, developing appropriate regulatory frameworks, and even minimizing negative social and economic effects. The creation of national strategies and a Latin American framework for AI implies facing up to the region's structural weaknesses and designing ways to exploit AI to boost its comparative advantages. It is also an opportunity for aligning objectives and values that are relevant to the spread of AI through Latin America as part of an intelligent approach to integration. These challenges demand forms of consensus that go beyond borders, based on inclusive regional and global governance models that focus AI development on achieving broad social benefits.

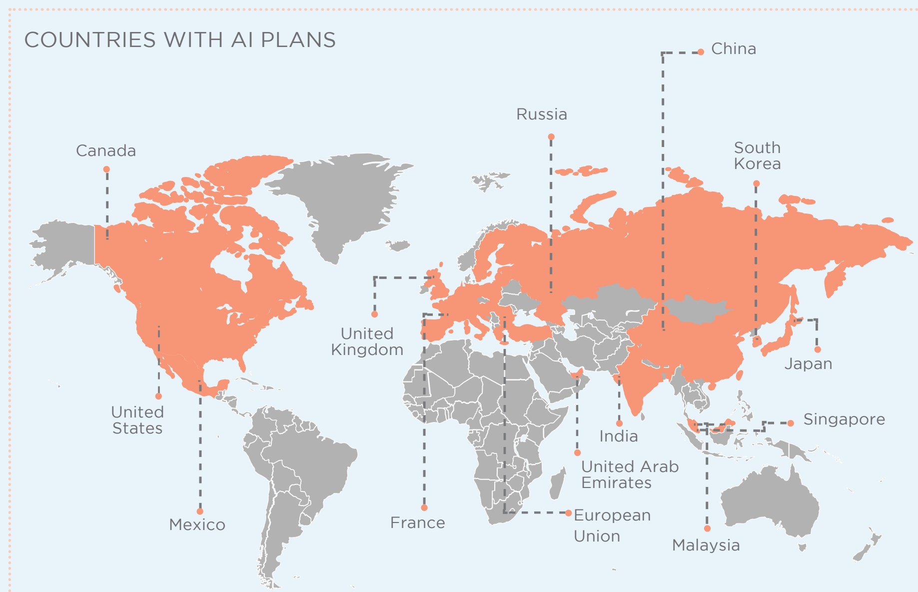

Predictive Integration

Artificial intelligence could prove to be an critical tool during trade negotiations, helping to accelerate the digital transformation of Latin American economies. The role of knowledge-based services and model prediction in data-driven macroeconomics.

The Productivity Leap

Armen Ovanessoff and Eduardo Plastino
Accenture Research

PRODUCTIVITY GAINS IN LATIN AMERICA HAVE BEEN MEDIOCRE EVEN DURING THE REGION'S PERIODS OF FASTER ECONOMIC EXPANSION. THIS ARTICLE EXAMINES HOW ARTIFICIAL INTELLIGENCE (AI) MAY BE TURNING THE PAGE IN THIS SENSE AND MAKING A SIGNIFICANT CONTRIBUTION TO THE GROWTH OF THE CONTINENT'S MAIN ECONOMIES.

One structural challenge for the global economy is the fact that the normal scenario in recent decades has been a slowing rate of growth (figure 1). Leading South American economies have been no exception. What is more, in recent decades, the region has seen more dynamic emerging markets in the Asia-Pacific region catch up with it fast, despite being previously less developed. 1

Moreover, the higher growth South America saw up to the beginning of the current decade has given way to a period characterized by the weakest regional performance since the so-called lost decade of the 1980s (figure 2). Key measures of economic efficiency are also trending sharply downward in the region, while labor-force growth is diminishing quickly and productivity is struggling to improve by much.

South America's recent poor performance puts a spotlight on its persistent productivity problems. In fact, productivity gains have been mediocre even during the region's periods of faster economic expansion. For example, during the 2000s, the group of five leading South American economies in our study improved their total factor productivity (TFP) by an annual average of only 0.1%. By comparison, South Korea's TFP grew by 0.7% and China's by 3.5% over the same period (figure 3).

During those good times, South American companies had the luxury of overlooking their productivity shortcomings because soaring revenues—especially from commodity exports and domestic consumption—ensured prosperity, even if their margins were squeezed. This was not sustainable. Today, the revenue drivers have fallen away, and the region's productivity problem is laid bare. A sustainable revival of growth must come hand-in-hand with productivity gains.

So where will new growth and productivity come from? Economists classify capital and labor as the traditional factors of production that drive expansion. Growth occurs when the stock of capital or labor increase or when they are used more productively.

In South America, the effectiveness in the use of capital has fallen for a decade. Moreover, the growth of the working age population is slowing down quickly, but labor is not becoming more productive quickly enough (figures 4 and 5).

Does this mean South America is experiencing the end of growth as we know it?

As grim as much of the data for the region—and indeed for most of the world—undoubtedly is, it misses an important part of the story. That missing element is how new technologies, especially AI, affect growth in the economy.

Economists have always thought of new technologies as driving growth through their ability to enhance TFP. This made sense for the technologies that we have seen until now: the great technological breakthroughs over the last two centuries boosted productivity dramatically.

Today, we are witnessing the take-off of another transformative set of technologies, commonly referred to as AI (see the box "What is artificial intelligence?"). Many see AI as being similar to past technological inventions. If we believe this, then we can expect some growth but nothing transformational.

However, AI has the potential to be not just another driver of TFP but an entirely new factor of production (figure 6). The key is to realize that AI is more than just another wave of technology. It is a unique hybrid of capital and labor.

Unlike previous technologies, AI creates an entirely new workforce. It can replicate labor activities at much greater scale and speed and even perform some tasks beyond the capabilities of humans. Furthermore, in some areas, it has the ability to learn faster than humans if not yet as deeply. For example, by using virtual assistants, 1,000 legal documents can be reviewed in a matter of days instead of taking three people six months to complete (Sobowale, 2016).

Similarly, AI can take the form of physical capital such as robots and intelligent machines. And unlike conventional capital, such as machines and buildings, it can actually improve over time, thanks to its self-learning capabilities.

Based on our analysis and modeling, we can illustrate what happens when AI is seen as a new factor of production rather than just a productivity enhancer. The impact on projected growth for Brazil, for example, is significant. As figure 7 illustrates, the first scenario is business-as-usual, assuming that AI has no effect. The second indicates the traditional view of AI as a TFP enhancer, where it has a limited impact on growth. The third scenario shows what happens when AI can act as a new factor of production: there is a marked effect on economic expansion. This ability of AI to complement and enhance traditional factors of production is where its true potential lies.

THREE KEY FACTORS ARE ENABLING AI GROWTH:

1. Rapid increases in computational power. Six years ago, in 2012, Google's Udi Manber and Peter Norvig (2012) reported that processing a single Google search query required roughly the same amount of computing power "as all the computing done—in flight and on the ground—for the entire Apollo program," which included 17 missions, taking Neil Armstrong and another 11 astronauts to the moon.

2. A huge fall in the cost of storing data, reaching a marginal cost of near zero. The cost of storing 1 gigabyte of data on a disk drive fell from US$185,000 in 1970 to US$277 in 1995 and to only US$0.02 last year.2

3. An explosion in digitized data As Barry Smyth, professor of computer

FIGURE 1

PANEL A. GLOBAL GDP GROWTH

A LONG-TERM DECELERATION

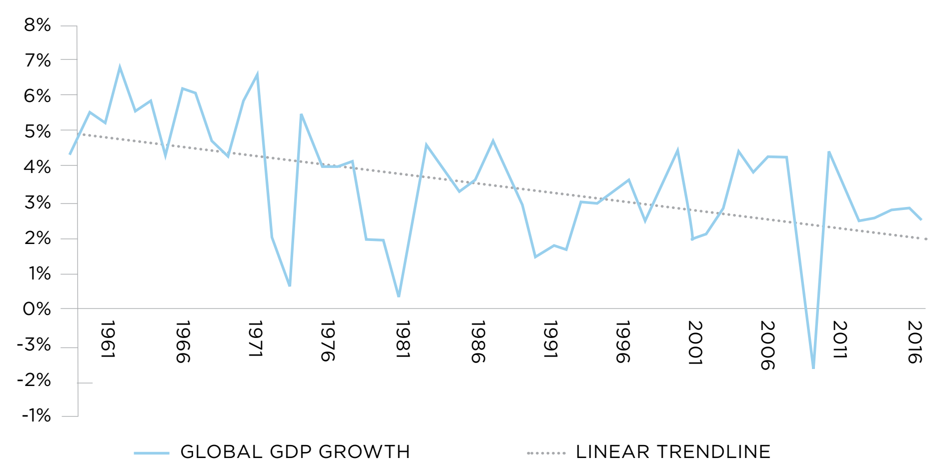

Source: World Bank

PANEL B. REGIONAL GDP GROWTH

MORE DYNAMIC EMERGING MARKETS IN ASIA HAVE LONG OUTPACED SOUTH AMERICAN ECONOMIES

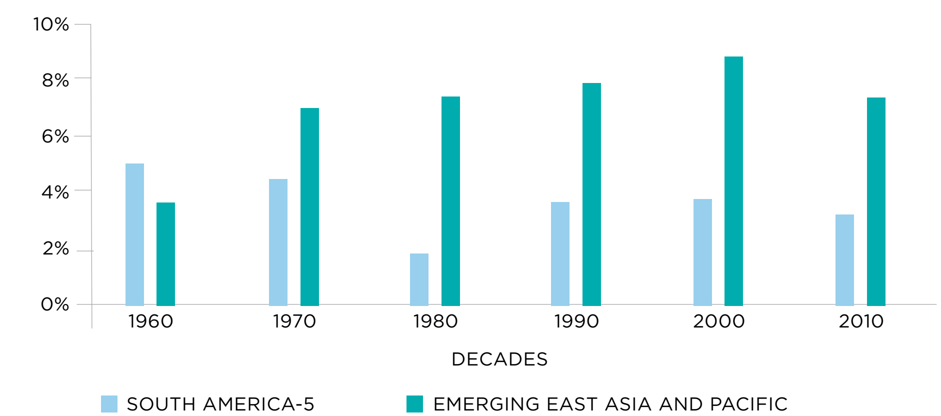

South America-5 is the unweighted average of Argentina, Brazil, Chile, Colombia, and Peru.
1960s = 1961–69; 1970s = 1970–79; 1980s = 1980–89; 1990s = 1990–99; 2000s= 2000–09; 2010s = 2010–16
Source: World Bank and Accenture (based on World Bank data)

SENSE, UNDERSTAND, ACT

AI is not a new field; much of its theoretical and technological underpinnings were developed over the past 70 years by computer scientists such as Alan Turing, Marvin Minsky, and John McCarthy. Today, the term refers to multiple technologies that can be combined in different ways to:

· Sense: Computer vision and audio processing, for example, are able to actively perceive the world around them by acquiring and processing images, sounds, and speech. The use of facial recognition at border control kiosks is one practical example of how it can improve productivity.

· Comprehend: Natural language processing and inference engines can enable AI systems to analyze and understand the information that is collected. This technology is used to power the language translation features of search engine results.

· Act: An AI system can take action through technologies such as expert systems and inference engines or undertake actions in the physical world. Auto-pilot features and assisted-braking capabilities in cars are examples of this.

All three capabilities are underpinned by the ability to learn from experience and adapt over time (figure 8). AI already exists to some degree in many industries but the extent to which it is becoming part of our daily lives is set to grow fast.

science at University College Dublin, told us: "data is to AI what food is to humans." So, in a more digital world, the exponential growth of data is constantly feeding improvements in AI.

THREE CHANNELS OF AI-LED GROWTH

Seen as the new factor of production, AI can drive growth in at least three important ways. First, it can create a new virtual workforce—what we call "intelligent automation." Second, AI can complement and enhance the skills and ability of existing workforces and physical capital. Third, like other previous technologies, AI can drive innovations in the economy, thus supporting TFP.

Intelligent Automation

The new AI-powered wave of intelligent automation is already creating growth through a set of features unlike those of traditional automation solutions.

The first feature is its ability to automate complex tasks from the physical world that require adaptability, agility, and learning. Consider the difficulties and dangers people face detecting perilous gases in a mine. Researchers from Peru's National Engineering University (UNI) have developed a four-wheeled robot that explores mines autonomously to detect methane, carbon dioxide, and ammonium. The robot uses sensors to detect its location and then generates improved routes and actions to be taken inside the mine as it collects information about gas levels and trends (Andina, 2016).

Whereas traditional automation technology is task-specific, the second distinct feature of AI-powered intelligent automation is its ability to solve problems across industries and jobs. Evidence of this is the widespread adoption of "chatbots" in customer service, virtual assistants that provide people with help in their native tongue. Today, these robots are eagerly employed by companies from Argentina's Banco Galicia to the Colombian airline Avianca and the Brazilian e-commerce platform ShopFácil.

The third and most powerful feature of intelligent automation is self-learning, enabled by repeatability at scale. Chilean startup The Not Company (or NotCo) developed an algorithm, nicknamed Giuseppe, that analyzes animal protein-based food products and generates recipes for vegan alternatives that not only have the same taste and texture but also offer better nutrition value. To do this, Giuseppe analyzes the molecular structure of foods and figures out similar structures based on combinations of vegan ingredients. Giuseppe makes "mayonnaise," for example, from potato starch, pea protein, and rosemary leaves. The larger its database grows, the more this algorithmic chef learns, and the more combinations it can produce (Baer, 2016). This self-learning aspect of AI is a fundamental leap forward. Whereas traditional automation capital degrades over time, intelligent automation assets can continually improve.

Labor and Capital Augmentation

A significant part of the economic growth from AI will come not from replacing existing labor and capital but from enabling them to be used much more effectively.

For example, AI can enable humans to focus on parts of their role that add the most value. Take a time-consuming and cumbersome process like recruitment, for example. The Chilean company AIRA (which stands

FIGURE 2
RECENT SOUTH AMERICAN PERFORMANCE
TOWARD A NEW LOST DECADE? (% CHANGE IN GDP)

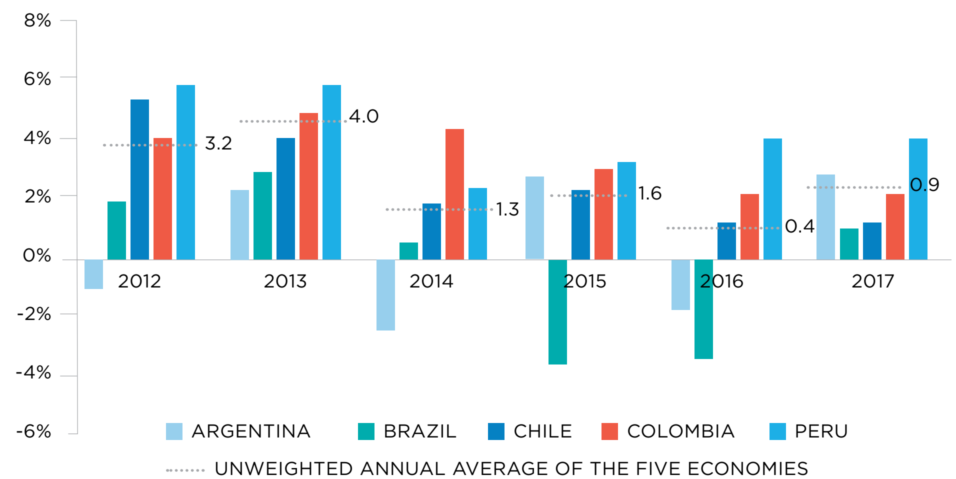

Source: IMF and Accenture (based on IMF data) for “artificial intelligence recruitment assistant”) has developed a system to publish vacancy announcements in the most widely used recruitment websites, read and rank all résumés, apply psychometric tests, and conduct video interviews with applicants. Applicants’ performance is assessed with emotion analytics, which translates their attention levels and facial expressions into numbers. At the end of this short process, human recruiters can focus their scarce time on conducting in-depth interviews with the best-qualified candidates (Pulsosocial, 2016).3

Moreover, AI augments labor by complementing human capabilities, offering employees new tools to enhance their natural intelligence. A number of companies in Brazil, for example, are preparing to incorporate hybrid intelligence systems into their postsales support services. This involves a robot collecting customer information from previous interactions with the firm such as product purchases, direct communication, or references over social media. It then provides the human attendant with information about the customer’s mood and any complaints and can also suggest promotions that might be relevant to each individual client.

AI can also improve capital efficiency, which is important for South America’s large industrial and manufacturing sectors. Take the case of Ubivis, a Brazilian startup established in 2014 with the ambition of helping manufacturing companies join the age of the internet of smart machines. Ubivis installs sensors and external drivers in existing industrial machines to collect large amounts of data about the client’s operations. Big data is then stored in the cloud and used as an input for machine-learning processes that make the client’s assets increasingly productive through, for example, predictive maintenance that solves problems before they become costly.

Innovation Diffusion

One of the least-discussed benefits of AI is its ability to drive new innovations as it spreads through the economy. Take driverless vehicles, probably the best-publicized AI product in development so far. As innovation begets further innovation, the impact of driverless vehicles on economies will eventually extend well beyond the automotive industry.

For example, the passenger—who is no longer driving—may well be engaged in mobile services, opening new opportunities for advertisers, retailers, media firms, and others to innovate new offerings. The insurance industry could generate more accurate risk assessments and new revenue streams from the masses of data that self-driving vehicles and their connected drivers produce. Public sector innovation opportunities also open up as real-time, accurate road and traffic data generated by vehicles and other sources open up new ways to charge for road usage and control congestion and pollution.

There could even be significant social benefits. Driverless vehicles are expected to reduce the number of road accidents and traffic fatalities dramatically, making the technology potentially one of the most transformative public health initiatives in human history. They could also give back independence to people who cannot drive due to disability, enabling them to take up jobs from which they were previously excluded. And, even among those who can drive, driverless cars will make traveling far more convenient, freeing up time that people can dedicate to work or leisure.

South America is already seeing driverless vehicles being used and designed for controlled environments, such as mines and ports, but as the

FIGURE 3
PANEL A. THE PRODUCTIVITY PROBLEM
LAGGING BEHIND DYNAMIC ASIAN ECONOMIES

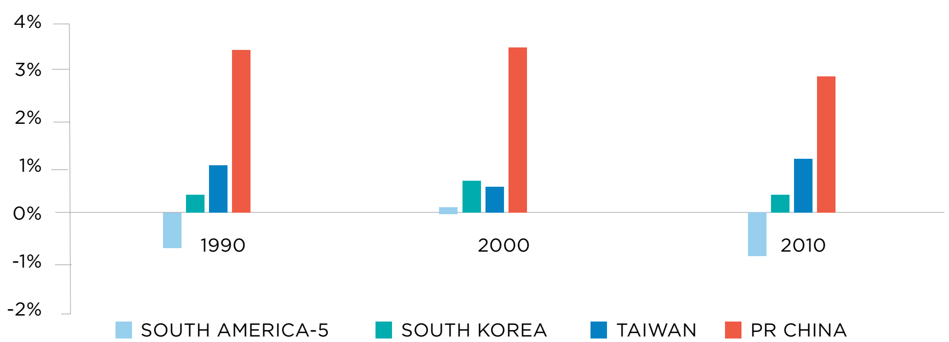

Total Factor Productivity (TFP) annual growth, average per decade.
TFP is the share of economic growth not explained by the contributions of labor or capital.
South America-5 is the unweighted average of Argentina, Brazil, Chile, Colombia, and Peru
Notes: 1990s = 1990–99; 2000s= 2000–09; 2010s = 2010–16.
“PR China” data is the official China gauge, as published by The Conference Board
Source: The Conference Board’s Total Economy Database and Accenture (based on the Conference Board’s Total Economy Database)

PANEL B. FALLING SOUTH AMERICAN PRODUCTIVITY
THE REGION’S ECONOMIES ARE BECOMING LESS PRODUCTIVE

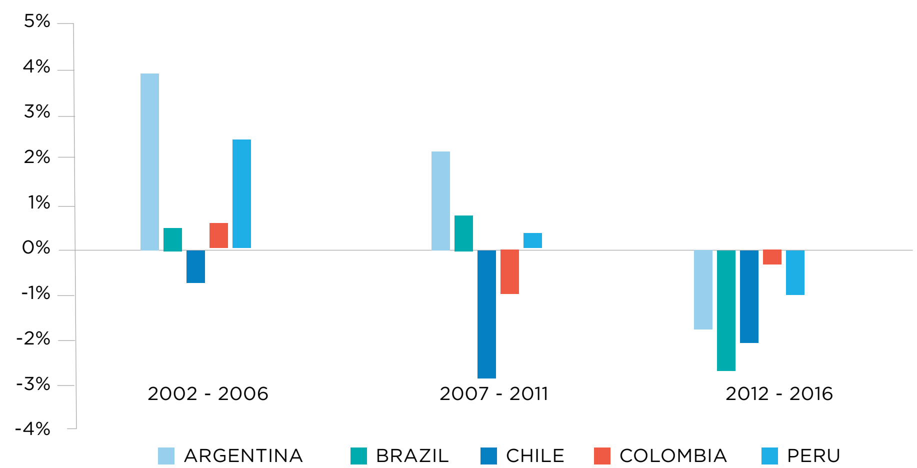

Average annual change in TFP for the period.
Source: Accenture (based on The Conference Board’s Total Economy Database)

FIGURE 4

CAPITAL EFFICIENCY

THE MARGINAL CAPITAL EFFICIENCY RATE, AN INDICATOR OF THE PRODUCTIVITY OF CAPITAL SUCH AS MACHINES AND BUILDINGS, HAS DROPPED SIGNIFICANTLY IN THE LAST DECADE

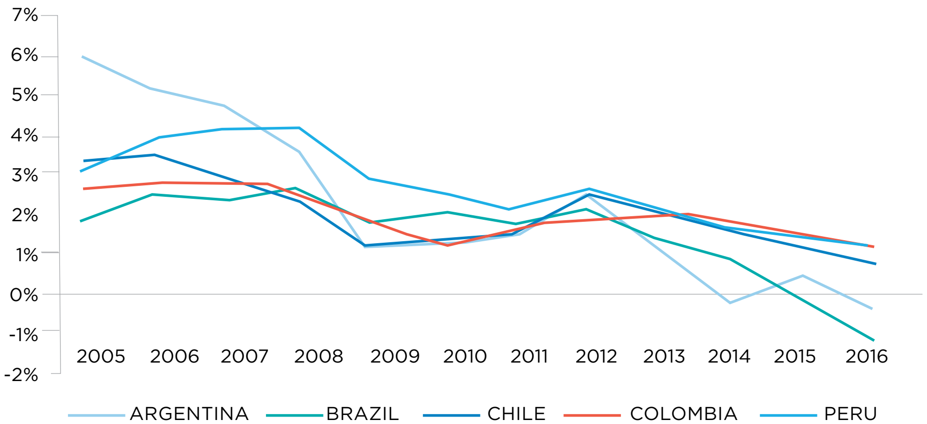

Marginal capital efficiency rate, 3-year moving average
Source: Accenture (based on World Bank data).

FIGURE 5

PANEL A. WORKING-AGE POPULATION GROWTH IS SLOWING QUICKLY

AS POPULATIONS AGE AND BIRTH RATES SLOW, FEWER PEOPLE ARE AVAILABLE TO PICK UP THE SLACK IN THE WORKFORCE

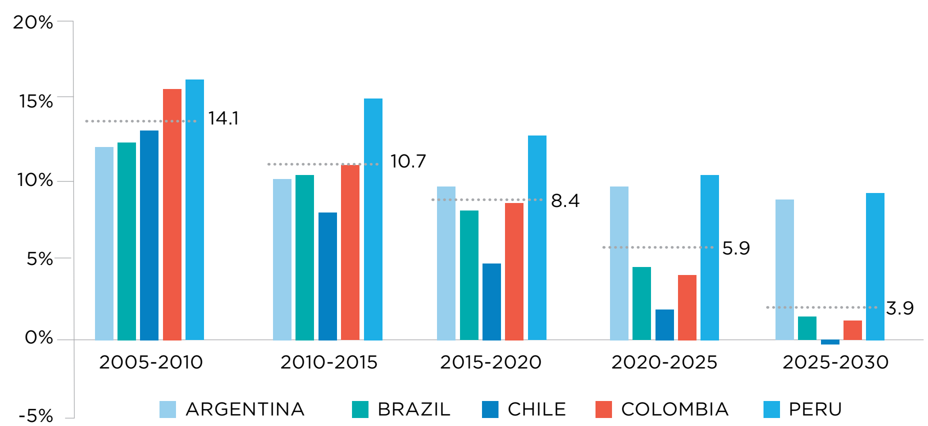

Note: Working-age population average growth per 1,000 residents. Working age defined as 15–59.
Source: UN Economic Commission for Latin America and the Caribbean (ECLAC) data and forecasts

technology and regulations advance, the opportunities will multiply.

A NEW FACTOR OF PRODUCTION

To understand the value of AI as a new factor of production, Accenture, in association with Frontier Economics, modeled the potential impact of AI on five economies that together generate some 85% of South America's economic output. Our results reveal notable opportunities for value creation. We find that AI has the potential to add up to 1% to annual economic growth rates across these countries—a powerful remedy for the low rates of recent years.

To estimate the economic potential of AI we compared two scenarios for each country. The first is the baseline, which shows the expected annual economic growth rate under current assumptions about the future. The second is the AI scenario, which shows expected economic growth once the impact of AI has been absorbed into the economy. As it takes time for the impact of a new technology to feed through, we used 2035 as the year of comparison (see appendix).

According to our modeling of Argentina, Brazil, Chile, Colombia, and Peru, AI would yield the highest economic benefit in absolute terms for Brazil, culminating in an additional US$432 billion in its gross value added (GVA) in 2035. This would mean a boost of 0.9 percentage points to growth that year. Chile and Peru could both increase their GVA—a close proxy for GDP—by an entire percentage point in 2035 thanks to AI. Colombia, meanwhile, could see an additional expansion of 0.8 percentage points.

Cross-country comparisons mask the significant impact that AI could have on seemingly lagging economies. For instance, AI is expected to raise Argentina's growth rate in 2035 from 3.0% to 3.6%. This implies the

PANEL B. LABOR PRODUCTIVITY GROWTH REMAINS POOR

LABOR OUTPUT PER HOUR WORKED IN 2017 US$ (CONVERTED TO 2017 PRICE LEVEL WITH UPDATED 2011 PPPS)

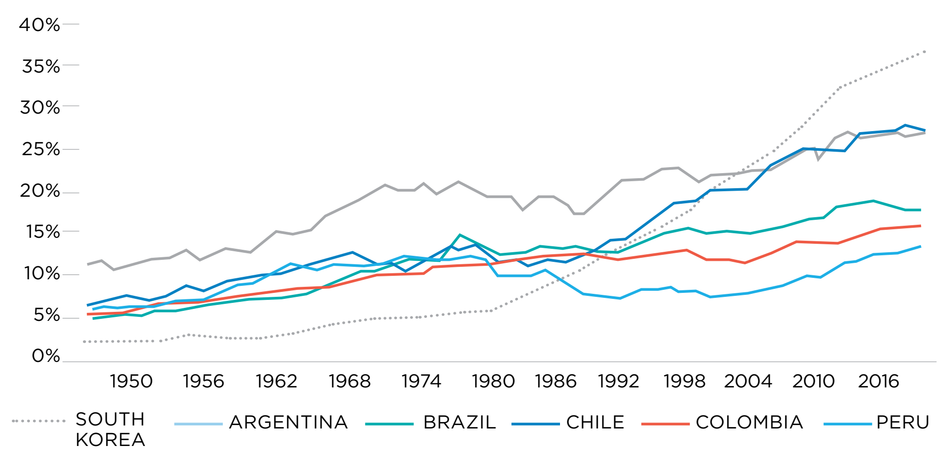

Source: The Conference Board's Total Economy Database lowest AI-driven boost among the five countries, but even this relatively modest contribution is still a sizeable amount: nearly US$59 billion of additional GVA, leading to a total GVA of US$702 billion in 2035.

Throughout South America, faster growth enabled by AI will reduce the number of years required for each national economy to double in size. Overall, AI is expected to unleash remarkable benefits across countries, redefining "the new normal" as a period of higher and longer-lasting economic growth.

POTENTIAL IMPACT

By focusing on individual countries, we can see the relative importance of the three channels through which AI has an effect. We compare the baseline-scenario size of each economy in 2035 with the AI scenario—that is, on in which AI has been absorbed into the economy.

The optimistic view about AI, in the words of futurist Ray Kurzweil (2014), is that it can help us make "major strides in addressing the grand challenges of humanity." In South America, forward-looking entrepreneurs such as Ubivis' Paulo Henrique Souza see it as a once-in-a-lifetime opportunity for the region to "leapfrog in technology."

AI, however, is not without problems. Entrepreneur Elon Musk has warned that it could become humanity's "biggest existential threat." The late British physicist Stephen Hawking feared that "the development of full artificial intelligence could spell the end of the human race" (Cellan-Jones, 2014). Even if these gloomy views do not come to pass, AI has the potential to increase unemployment and inequality.

So, is it a bad or a good thing?

The truth is, it all depends on how we manage the transition to an era of AI. To fulfill the promise of AI as a new factor of production that could reignite economic growth and benefit all, relevant stakeholders must be thoroughly prepared—intellectually, technologically, politically, ethically, and socially—to address the challenges that arise as AI becomes more integrated into our lives. Five action areas are critical.

Preparing the Next Generation

A recurrent theme in interviews we conducted with businesspeople and academics from across South America was the region's shallow talent pool. The International Labour Organization calculates that only about 20% of

FIGURE 6
THE AI GROWTH MODEL
OUR MODEL ADAPTS THE TRADITIONAL GROWTH MODEL BY INCLUDING AI AS A FACTOR OF PRODUCTION

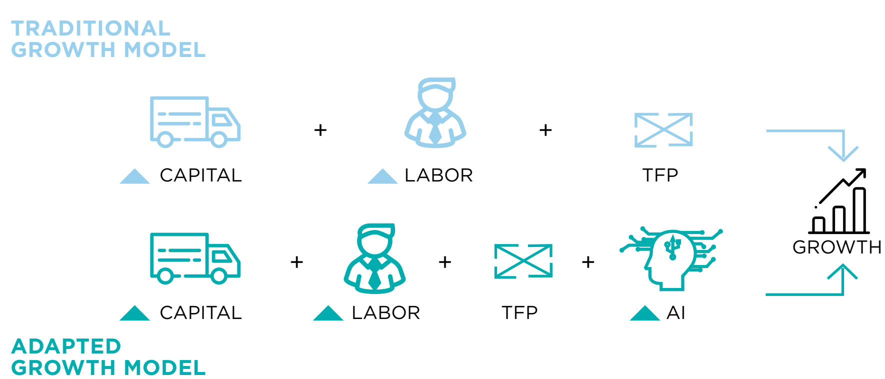

Note: ▲ Indicates the change in that factor.
Source: Accenture.

FIGURE 7
THREE GROWTH SCENARIOS FOR BRAZIL'S ECONOMY
AI AS A NEW FACTOR OF PRODUCTION CAN LEAD TO SIGNIFICANT GROWTH OPPORTUNITIES FOR BRAZIL'S ECONOMY

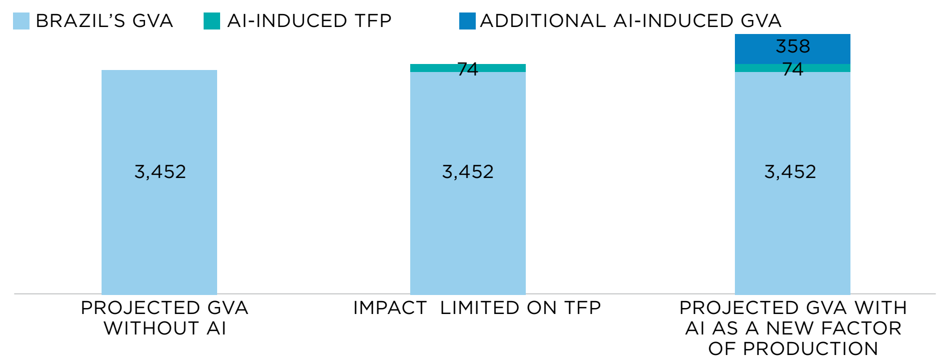

Note: Brazil's gross value added (GVA) in 2035 (billions of US$).
Source: Accenture and Frontier Economics.

FIGURE 8
EMERGING AI TECHNOLOGIES

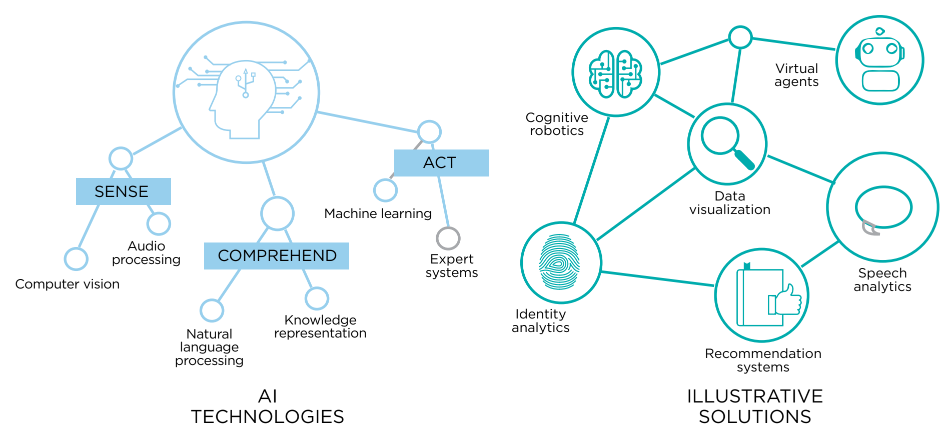

Source: Accenture.

South American workers are in jobs that require high-level skills, compared with over 40% in the European Union and the United States (figure 12).4

To address this, improving the quality of national education systems will be crucial, as will be increasing access to tertiary education in different countries. But preparing the next generation will not suffice. The current generation of South American workers will also need support to adapt to the AI economy.

This will take place against a backdrop in which, as AI eliminates the need for humans to perform a number of tasks, it frees up time for people to learn about areas where they can add more value. In this context, technical skills will be required to design and implement AI systems, exploiting expertise in many specialties including robotics, computer vision, and pattern recognition. In this area, the region has work to do. South American countries were in the bottom half of the 70 economies included in the latest OECD science tests applied to 15-year-olds in 2015 (OECD, 2018).

But interpersonal skills, creativity, and emotional intelligence will become even more important than they are today. Perhaps, South America's widely recognized culture of openness and ease with relationship building and communication could prove an advantage here (see the box "Building Skills for the Age of Intelligent Machines").

Strengthening AI Ecosystems

Innovation flourishes when relations between startups, large companies, academic researchers, government agencies, and other key stakeholders are regular and intense. Unfortunately, innovation ecosystems in South America tend to be weak. Previous

FIGURE 9

THE ECONOMIC IMPACT OF AI

AI HAS THE POTENTIAL TO INCREASE ANNUAL ECONOMIC GROWTH RATES IN SOUTH AMERICA BY UP TO 1 PERCENTAGE POINT IN TERMS OF GROSS VALUE ADDED

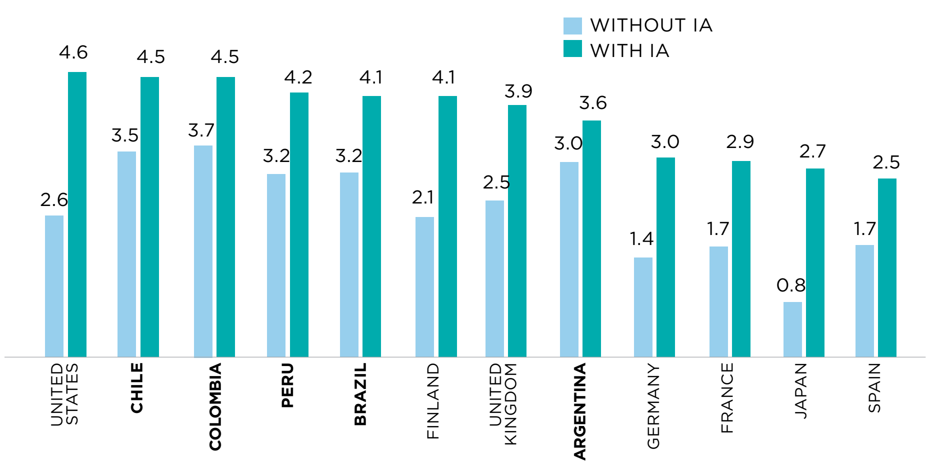

Source: Accenture and Frontier Economics.

FIGURE 10

TIME FOR ECONOMIES TO DOUBLE IN SIZE

AI PAVES THE WAY TO FASTER ECONOMIC GROWTH

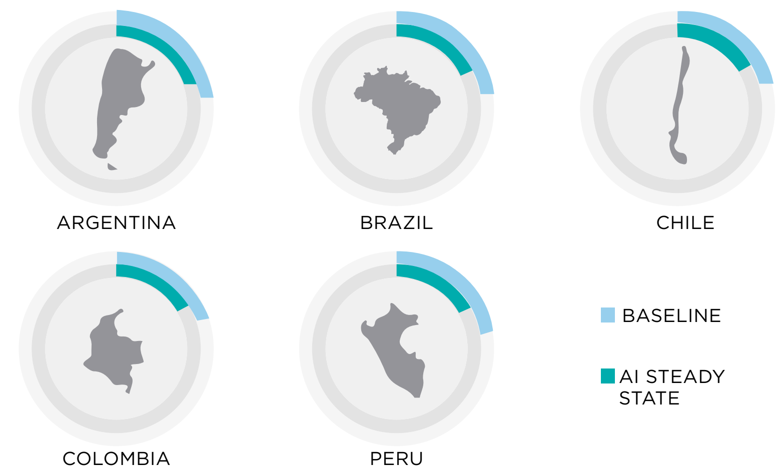

Number of years for the economy to double in size (a full circle represents 100 years).
Source: Accenture and Frontier Economics.

Accenture research highlights how low trust levels in the region contribute to weak levels of collaborative innovation (Ovanessoff, Plastino, and Faleiro, 2015). "Global trends point to collaboration and cocreation, but in our country [lack of] trust is tripping us up," laments Renzo Pruzzo, general manager of Chile's Innovation Club, a cross-industry organization.

Ana Maguitman, a researcher at Argentina's National University of the South and Argentina's National Scientific and Technical Research Council, acknowledges the need to build trusting relationships across institutions, not only between businesses. For her, the penny dropped after the liaison office at her university identified her work as having potential for commercial development: "They educated me about technology transfer and how to communicate with companies. This is still something new for us."

Some far-sighted firms in the region are already exploring the opportunities offered by global ecosystems. Ubivis, the Brazilian startup, is a keen user of open-source IoT computer programs offered by organizations such as the Apache Software Foundation and Eclipse. As the company's CEO explained to us: "We build on the work of these global organizations. We add the vital 30% on top of their 70%."

Encourage Suitable Regulation

As autonomous machines take over tasks that have exclusively been undertaken by humans, current laws will need to be revisited. For instance, the State of New York's 1967 law that requires drivers to keep one hand on the wheel was designed to improve safety but may inhibit the uptake of semi-autonomous safety features, such as automatic lane centralization (Kessler, 2015).

In other cases, new regulations are called for, and delays in implementing new laws may undermine South America's progress in the adoption and development of AI technologies. In Brazil, for instance, science-related legislation tends to be "reactive," says Yasodara Córdova, a Brazilian Research Fellow at the Digital Harvard Kennedy School of Governance. "This has two somewhat contradictory consequences, both negative. It results in legal uncertainty, which discourages investment. At the same time, those seeking to exploit ethically questionable uses of new technology can profit from the legal vacuum," she notes. Likewise, Colombian IT lawyer Natalia Ospina Díaz (2016) explains that her country isn't "legally prepared" for today's pace of technological change.

Here too, AI itself can be part of the solution, creating adaptive, self-improving regulation that closes the gap between the speed of technological evolution and the regulatory response to it. In the same way that intelligent solutions combined with massive data can guide decision-making in areas such as urban planning, healthcare, and social services, they could also be used to update regulation in light of new cost-benefit evaluations.

Advocate a Code of Ethics

Intelligent systems are rapidly moving into social environments that were once only occupied by humans.

This is opening up ethical and societal issues that can slow down the progress of AI. These range from how to respond to racially biased algorithms to whether autonomous cars should give preference to their driver's life over that of others in the case of an accident. Given how prevalent intelligent systems will be in the future, policymakers need to ensure the development of a code of ethics for the AI ecosystem.

Ethical debates need to be supplemented by more tangible standards and best practices in the development of intelligent machines. As a segment of AI, the robotics industry is already ahead in setting universal standards for its operations. Business standards regarding robots produced by the British Standards Institution (BSI) are a step in the right direction.

FIGURE 11
AI'S IMPACT ON NATIONAL ECONOMIES

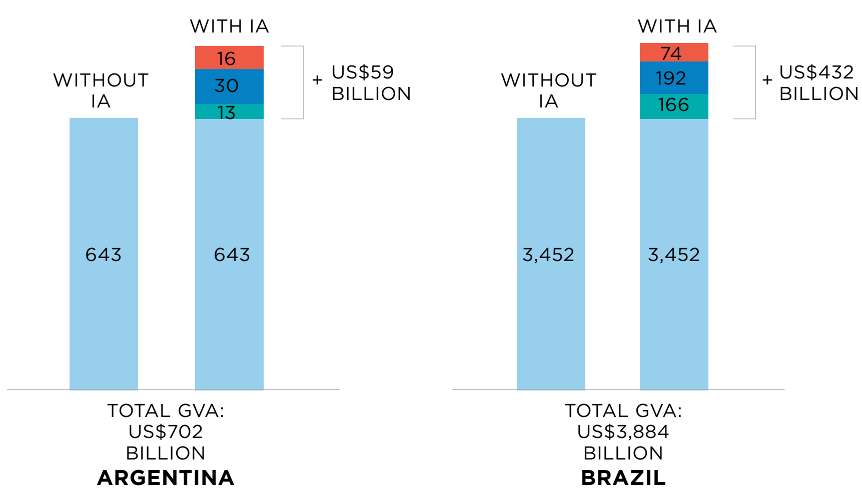

Source: Compiled by the authors.

Minimize Risks to Social Cohesion

A widespread and legitimate concern of many commentators is that AI will eliminate jobs, worsen inequality, and erode incomes. Given the region's extremely high inequality level—with 10% of the population already controlling around 70% of the wealth in the wider Latin America5 (Bárcena Ibarra and Byanyima, 2016)—this risk must be taken extremely seriously, and the region must prepare to face it.

The prospect of job losses driven by breakneck technological progress is the main reason why some places in countries from Canada to the Netherlands have begun testing pilot universal basic income schemes (Tencer, 2016; Brown Hamilton, 2016). 6 "The need for a universal basic income will become increasingly clear," warns professor Guillermo Simari, chair of the Artificial Intelligence R&D Lab at Argentina's National University of the South.

But income is only part of the equation. We may see shifts in the value society ascribes to the roles and contributions of people, machines, and communities. How will we treat paid work versus unpaid work? Will we tax robots? Will sections of society feel freed or stripped of their dignity and self-worth if paid work is no longer an option for them? Such questions about the structure of society and social contracts must be carefully examined as we plan this journey.

Yet at the same time, policymakers must also articulate the common benefits that AI offers to society at large. For example, large sectors of the workforce will benefit from more stimulating work and greater job satisfaction. An Accenture survey highlighted that 84% of managers across 14 countries believe AI will make them more effective and their work more interesting.

Beyond the workplace, AI promises to alleviate some of the world's greatest problems, such as climate change (through more efficient transportation) and poor access to healthcare (by reducing the strain on overloaded systems). Benefits like these should be clearly explained to encourage a more complete and realistic outlook on AI's potential.

AI offers a rare opportunity for South American economies to address their productivity deficit and increase their dynamism on a more sustainable basis.

The good news is that AI is already becoming a reality for many sectors across the region, and the appetite from business, government

70%

OF LAC'S WEALTH IS CONCENTRATED IN THE HANDS OF JUST 10% OF ITS POPULATION

FIGURE 12
SHARE OF THE WORKFORCE IN JOBS THAT REQUIRE HIGH-LEVEL SKILLS

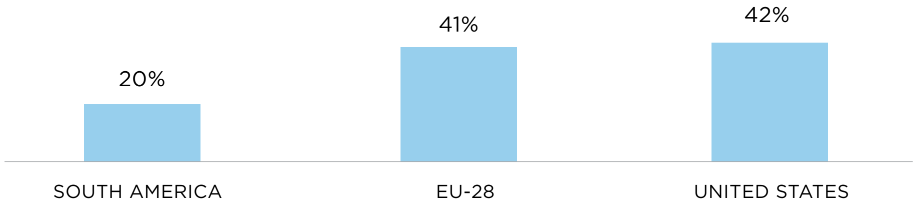

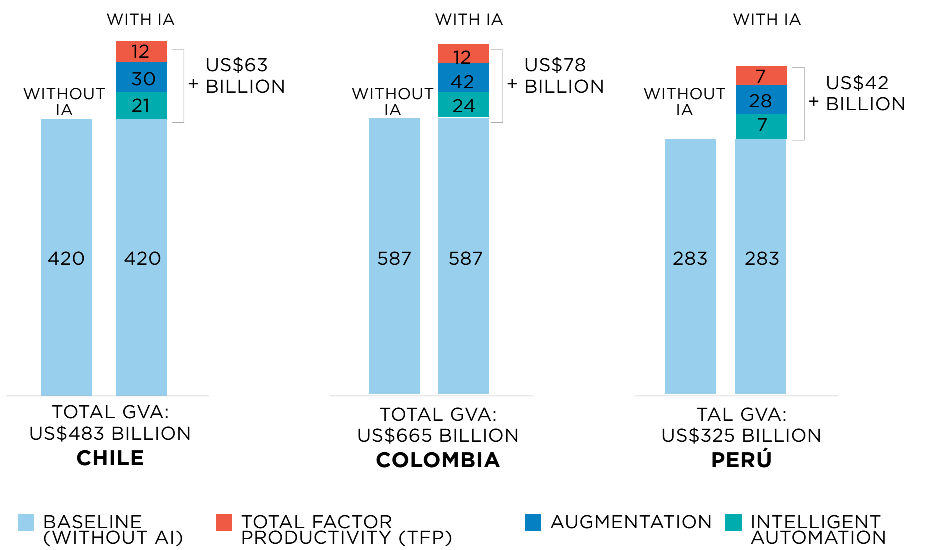

Source: International Labour Organization (ILO) estimates for 2018, employment by occupation

and individuals appears no less in South America than in the most technologically advanced parts of the world. Indeed, groundbreaking AI systems and applications are being designed and built within the region.

At the same time, the region must make fundamental improvements to some basic areas, such as education systems and research and innovation ecosystems, in order to capture the broad and deep benefits that AI promises.

That said, the largest challenges to capitalizing upon the opportunity that AI represents are no different in South America than they are anywhere else in the world. South American business leaders and policymakers should not think of themselves as "catching up"—as they are often used to feeling—when it comes to AI. Rather, they should actively engage with their peers around the world to steer AI toward a productive and sustainable source of social and economic growth for all.

DEVELOPING CAPACITIES

In recent research on the future of labor markets in Latin America, Accenture found that almost one in three workers employed in the formal economy in most of the region's largest markets spend over 75% of their time on routine tasks. Since those are the tasks that intelligent machines—computers or robots—are more likely to perform, such workers are more at risk of having their jobs automated away or transformed in the coming years.

FIGURE 13
PROBABILITY OF AUTOMATION BY 2020, BASED ON ROUTINE TASK CONTENT

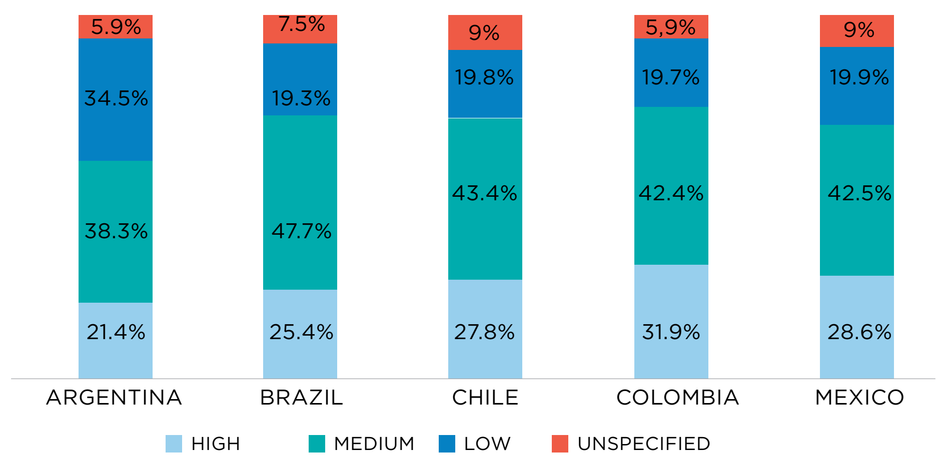

Source: Accenture.

This does not mean that job losses are inevitable. Just because jobs can be automated, it does not mean that they will be. We have also found that technology can help such workers obtain the skills they will need for the jobs that will be created, through means that vary from online courses to training programs that use virtual and augmented reality. Moreover, employers can reconfigure jobs after some tasks are automated, keeping existing workers in a changed function. Nevertheless, Latin American countries cannot escape the need to upgrade the skills of vulnerable workers and make sure that young people entering the labor force have the skills they need. This is crucial to avoid more informality—which already affected 46.8% of nonagricultural workers in 2015, according to the International Labour Organization—falling wages for many, and a reversal of the region's recent progress in reducing poverty and inequality.
Our suggestions of concrete steps that businesses, their partners in the NGO

community, and state training institutions can take to guide a responsible digital transformation include the following: 1) assessing the economic and social trends that will create new jobs in the coming years in Latin America; 2) focusing on building the enduring human skills that will be increasingly relevant in the job market, as well as new technical skills, and 3) using new teaching methods, including experiential learning and technology-based training7.

NOTES

[1]The authors would like to thank Mark Purdy and Roberto Frossard for their support and input.
[2]According to IDC research cited in the following article by Computerworld's senior reporter, Lucas Mearian (2017).
[3]http://site.airavirtual.com/.
[4]International Labour Organization, "Employment by Occupation—ILO Modelled Estimates, Nov. 2017," available at: www.ilo.org/ilostat.
[5]2014 figures for Latin America.
[6]For more on a recently aborted pilot program granting universal basic income to unemployed people in Finland, see, for example, Goodman (2018) and Jauhiainen and Mäkinen (2018).
[7]For a more in-depth analysis, see Plastino, Zuppolini, and Govier (2018).
[8]For further details on the importance of national absorptive capacity, see Purdy and Davarzani (2015).

REFERENCES

Andina 2016. "Robot minero también tendría utilidad en agricultura y desastres naturales." Agencia Peruana de Noticias. July 13.
Baer, D. 2016. "This Startup Is Using Machine Learning to Create Animal Product Substitutes." *Business Insider*. April 26.
Bárcena Ibarra, A., and Byanyima, W. 2016. "Latin America Is the World's Most Unequal Region. Here's How to Fix It." WEF. January 17.
Brown Hamilton, T. 2016. "The Netherlands' Upcoming Money-for-Nothing Experiment." *The Atlantic*. June 21.
Cellan-Jones, R. 2014. "Stephen Hawking Warns Artificial Intelligence Could End Mankind." BBC News. December 2.
Frey, C.B., and Osborne, M.A. 2013. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" Working Paper. Oxford Martin School and University of Oxford.
Goodman, P.S. 2018. "Finland Has Second Thoughts About Giving Free Money to Jobless People." *The New York Times*. April 24.
Jauhiainen, A., and Mäkinen, J.H. "Universal Basic Income Didn't Fail in Finland. Finland Failed It." *The New York Times*. May 2.
Kessler, A.M. 2015. "Law Left behind as Hands-Free Cars Cruise." *Star Tribune*. May 3.
Kolbjørnsrud, V., Amico, R., and Thomas, R.J. 2016. *The Promise of Artificial Intelligence: Redefining Management in the Workforce of the Future*. Accenture.
Kurzweil, R. 2014. "Don't Fear Artificial Intelligence." *Time*. December 19.
Manber, U., and Norvig, P. 2012. "The Power of the Apollo Missions in a Single Google Search." Google Search Blog. August 28.
Mearian, L. 2017. "CW@50: Data Storage Goes from $1M to 2 Cents per Gigabyte (+Video)." *Computerworld*. March 23.
OECD. 2018. *PISA 2015 Results in Focus*. OECD.
Ospina Díaz, N. 2016. "¿Está Colombia preparada jurídicamente para enfrentar el desafío tecnológico?" Portafolio Blog. June 1.
Ovanessoff, A., Plastino, E., and Faleiro, F. 2015. *Why Brazil Must Learn to Trust in Collaborative Innovation*. Accenture.
Plastino, E., Zuppolini, M., and Govier, M. 2018. *Building Latin America's Skills For The Age of Intelligent Machines*. Accenture.
Pruzzo, R. 2015. "La confianza como pieza clave para la co-creación." *Diario Financiero*. September 13.
Pulsosocial, 2016. "Inteligencia artificial 'made in Chile': AIRA acelera reclutamiento y selección de personal." *Pulsosocial*. October 20.
Purdy, M., and Davarzani, L. 2015. *The Growth Game-Changer: How the Industrial Internet of Things Can Drive Progress and Prosperity*. Accenture Strategy.
Sobowale, J. 2016. "How Artificial Intelligence Is Transforming the Legal Profession." *ABA Journal*. April.
Tencer, D. 2016. "Basic Income Coming To P.E.I.? Legislature Passes Motion Unanimously." *The Huffington Post Canada*. December 7.

APPENDIX: MODELING THE IMPACT OF AI ON GVA

AI has the potential to unleash a broad-based disruptive impact on society, creating a variety of economic benefits. While some of these benefits can be measured, others, such as consumer convenience and time savings, are far more intangible in nature. Our analysis focuses on measuring the impact of AI on GVA.

We began with a modified growth model developed by Robin Hanson, professor of economics at George Mason University, Virginia, United States. We looked at the additional increase in growth that would occur as a result of AI by contrasting it with the baseline growth rate.

In our model, we defined labor as a continuum of tasks that can either be performed by a human or artificial intelligence—not work solely done by humans. The intent was to introduce intelligent systems as an additional workforce capable of handling activities that require an advanced level of cognitive agility.

To estimate the shares of workers' tasks that could be performed by intelligent machines (AI absorption rates), we drew on research by Frey and Osborne (2013), who take a task-based approach to identifying roles and occupations that are affected by AI. The estimates are aggregated at country-level, taking into account the different mix of occupations and industries within each country.

These figures were adjusted to reflect the following:

- We assume that employment will be constant in the long term.
- Differences between AI's technological potential and the actual potential achieved: we considered the uptake of AI from zero to the maximum technological potential. We assumed that a 50% uptake would be reasonable in the period analyzed, that is, AI substitution is assumed to achieve 50% of its technological potential.
- A key driver of the impact of AI on growth is how well each country is positioned to benefit from the emergence of new technologies and how ready it is to integrate them into its economy, as measured by what we refer to as a country's "national absorptive capacity."1

With these calculations and adjustments, we arrived at our final estimates of the AI absorption rates used in our macro model. Along with the quantitative model, we supplemented our research by conducting interviews with businesspeople and experts from a range of different disciplines and secondary research to give insight into the capacity of AI to generate economic growth.

A New Phase in Globalization

Anand S. Rao
PriceWaterhouseCoopers

THE FOURTH INDUSTRIAL REVOLUTION MARKS THE START OF A NEW STAGE IN INTERNATIONAL ECONOMIC RELATIONS, PARTICULARLY THROUGH THE SPREAD OF APPLICATIONS OF ARTIFICIAL INTELLIGENCE (AI) TO BUSINESS. OVER THE NEXT DECADE, PROFOUND GLOBAL DISRUPTIONS ARE EXPECTED IN DIFFERENT ECONOMIC SECTORS. THIS ARTICLE ANALYZES THE OPPORTUNITIES THAT THESE TRANSFORMATIONS MAY BRING, AND THE RISKS ASSOCIATED WITH THEM.

Artificial intelligence (AI) is the simulation of human intelligence processes by machines, especially computer systems. AI includes the theory and development of computer systems that perform tasks that normally require human intelligence, such as speech recognition, decision-making, visual perception, and others. AI is becoming a ubiquitous form of intelligence that can see, hear, speak, smell, feel, understand gestures, interface with your brain, and dream. One application of AI, machine learning (ML), is the development of computer programs that can access data to learn for themselves without being programmed with explicit rules.

AI research has recently been described as the “fourth Industrial Revolution” (Schwab, 2016; Kelnar, 2016). Innovative, groundbreaking technologies are connecting millions of people and machines for the purpose of automating and completing tasks that were once impossible. The growing sophistication of AI and ML technologies is transforming our lives both as consumers and within enterprise.

But why has AI suddenly caught fire when research in this field began in the 50? AI is growing at a rapid rate in recent years as a result of the abundance of data, accelerating improvements in processing power, easy and cheap access to these through cloud computing, improvement in algorithms, and easy access to open-source tools and techniques.

FACTORS BEHIND THE BOOM

Today, AI is in use in our daily lives and has reached a historical moment because of six converging factors:

Big data: Computers have given us access to vast amounts of data, both structured (in databases and spreadsheets) and unstructured (such as text, audio, video, and images). All of this data documents our lives and improves humans’ understanding of the world. As trillions of sensors are deployed in appliances, packages, clothing, autonomous vehicles, and elsewhere, big data will only get bigger. AI-assisted processing of this information allows us to use this data to discover historical patterns and make more efficient predictions and more effective recommendations, among other things. As data generation and storage have become the norm for businesses, data availability has grown exponentially.

Processing power: Accelerating technologies such as cloud computing and graphics processing units (GPUs) have made it cheaper and faster to handle large volumes of data with complex AI-empowered systems through parallel processing. Innovation in processing power has allowed for more scientific experimentation in the field of AI. GPUs are specialized electronic circuits that minimize the time needed for complex ML algorithms to train by speeding up calculations with matrix multiplications. As a result, algorithms that train neural

54%
OF THE WORLD'S EXECUTIVES ARE INVESTING IN AI

networks, which rely heavily on large amounts of training data and thousands of parameters, have vastly improved. Researchers can more efficiently experiment with training these networks for complex tasks, such as image recognition. In the future, deep learning chips—a key focus of research today—will push parallel computation further.

A connected world: social media platforms have fundamentally changed how individuals interact. This increased connectivity has accelerated the spread of information and encouraged the sharing of knowledge, leading to the emergence of a collective intelligence, including open-source communities that develop AI tools and share applications.

Open-source software and data are accelerating the democratization and use of AI, as can be seen in the popularity of open-source ML standards and platforms such as TensorFlow, Caffe2, PyTorch, and Parl.ai. An open-source approach can mean less time spent on routine coding, industry standardization, and wider application of emerging AI tools.

Researchers have made advances in several aspects of AI, particularly in "deep learning," which involves layers of neural networks designed in a fashion inspired by the human brain's approach to processing information. Deep learning neural networks (multilayer perceptrons) were first envisioned in 1965 but are only now being built and tested. We have been able to achieve human-level image and speech recognition with these advancements. The error rates have dropped from 25% to 5% in a matter of a few years. Another emerging area of research is "deep reinforcement," in which the AI agent learns with little or no initial input data, by trial and error. Similar achievements will continue to build across the field, and eventually AI will impact everything we touch, from manufacturing to social experiences.

Accelerating returns: Competitive pressures have fueled the rise of AI, as businesses have used algorithms and open-source software to boost their competitive advantage and augment their returns through, for example, increasing personalization of consumer products or utilizing intelligent automation to increase their productivity.

The convergence of these factors has helped AI move from in vitro (in research labs) to in vivo (in everyday lives). Established corporations and start-ups alike can now pioneer AI advances and applications. Indeed, many people are already using AI-infused systems, whether they realize it or not, to navigate cities, shop online, find entertainment recommendations, filter out unwanted emails, or share a journey to work. AI is already here, then, and many corporate executives perceive its potential value. In a 2017 PwC survey of global executives, 54% reported making substantial investments in AI, while a lack of digital skills remains an important concern. As organizations continue to invest in tools, data optimization, people, and AI-enabled innovations, the realized values are expected to take off: growing from US$1.4 billion in annual revenue from AI-enabled systems in 2016 to US$59.8 billion by 2025, according to one research study.

LAND OF OPPORTUNITY

The analysis carried out by PwC (Rao, Verweij, and Cameron, 2017; Gilham et. al., 2018) gauges the economic potential for AI between now and 2030 and includes regional economies and eight global commercial sectors. Through our AI Impact Index, we also look at how improvements to personalization/customization, quality, and functionality could boost value, choice, and demand across nearly 300 use cases of AI, along with how quickly transformation and disruption are likely to take hold. Other key elements of the research include in-depth sector-by-sector analyses (Gilham et. al., 2018).

What emerges clearly from the analysis we carried out for this report is just how big a game changer AI is likely to be and how much potential value is up for grabs. AI could contribute up to US$15.7 trillion (or an additional 14%) to the global economy in 2030, more than the current output of China and India combined. Of this, US$6.6 (5.8%) trillion is likely to come from increased productivity, and US$9.1 trillion (8%) is likely to come from consumption effects. While some markets, sectors, and individual businesses are more advanced than others, AI is still at a very early stage of development overall. From a macroeconomic point of view, there are therefore opportunities for emerging markets to leapfrog more developed counterparts. Within the business sector, one of today's start-ups or a business that hasn't even been founded yet could be the market leader in ten years' time.

Regarding product enhancement impacts, nearly all the GDP impact is expected to derive from increases in varieties of goods as well as increases in the quality of these, with only a negligible impact from increases in time savings for consumers using AI-enhanced products. This not only reflects the fact that the base (%) increase in time savings is relatively small on average over the year (per consumer) but also that this increase in labor supply (availability) is not significant enough to actually incentivize an increase in the supply of labor meaningfully.

Notably, the consumption-side impacts are more delayed but are surprisingly large in nature, overtaking the labor productivity contribution to GDP gains in the late 2020s. Both the delayed and large nature of these impacts can be explained by the complex (initial) transmission mechanism from these product enhancements to consumption.

TABLE 1
RISKS POSED BY ARTIFICIAL INTELLIGENCE

1. PERFORMANCE RISKS	2. SECURITY RISKS	3. CONTROL RISKS	4. ETHICAL RISKS	5. ECONOMIC RISKS	6. SOCIETAL RISKS
Risk of errors Risk of bias Risk of opaqueness or black box risk Risk of explainability Risk of stability of performance	Cyberintrusion risks Privacy risks Open-source software risks	Risk of AI going "rogue" Inability to control malevolent AI	"Lack of values" risk Value alignment risk Goal alignment risk	Job-displacement risks "Winner-takes-all" concentration of power risk Liability risk Reputation risk	Risk of autonomous weapons proliferation Risk of intelligence divide

Source: PwC Data & Analytics.

14%+
THE BOOST AI COULD GIVE THE GLOBAL ECONOMY

Since the mechanism for transmitting product enhancements on the consumption side—particularly the entry of new firms—takes considerable time, the effect of product enhancements on GDP takes substantially longer to appear than effects driven by productivity. However, the impact on GDP is substantial once the transmission has taken place, which predominantly reflects the size of the increase in the affordability of goods due to new firm entry (figure 1).

However, there are two important caveats which must be considered when interpreting our analysis.

Our results estimate the upwards pressure on GDP as a result of AI only, under the ceteris paribus assumption. Our results may not be directly reflected in future economic growth figures, as there will be many positive or negative forces that either amplify or cancel out the potential effects of AI (e.g., shifts in global trade policy, financial booms and busts, major commodity price changes, geopolitical shocks, and so on).

As mentioned previously, our economic model results are compared to a baseline of long-term steady-state economic growth. The baseline is constructed from three key elements: population growth, growth in the capital stock, and technological change. The assumed baseline rate of technological change is based on average historical trends. Therefore, since AI has already been introduced prior to the starting evaluation period of this study, the component of these forecasts driven by technological change will already have factored into past trends in AI's impact on GDP. As a result, it is difficult to quantify the exact fraction of AI's impact on GDP that will be additional to historical average growth rates (i.e., additional to the baseline forecast).

However, our study is specifically focused on the AI technologies that are yet to be implemented and are expected to be implemented between 2017 and 2030. As a result, an underlying but reasonable assumption we make here is that the scale and impact of these AI technologies will be above the current trend in AI's uptake and impact. Based on this premise, our study is centered on estimating the total marginal economic impact of yet-to-be-implemented AI between 2017 and 2030—not including the AI that has already been implemented prior to this study (which is implicitly included in the baseline growth assumption). This also means that, while our study results imply that average economic growth rates between 2017 and 2030 will be raised due to AI's impact, we do not make claims outside of this time interval. As a result, we do not interpret that AI will impact the fundamental long-term growth rate of the global economy.

These two factors mean that our results should be interpreted as the potential size of the economic prize associated with AI over the period of our study, as opposed to direct estimates of the impact of AI on long-run economic growth.

One other economic caveat which must be addressed is that it is assumed that firms can enter and compete in the future under circumstances that are identical to today's. This assumes that data ownership as a mechanism for doing business will be regulated under the same antitrust considerations that the pre-AI economy enforced. However, if firms maintain almost exclusive ownership of data, or are able to build a sufficiently large moat around this, increased dynamic entry and competition may be less feasible.

All geographic regions of the global economy will experience economic benefits from AI. North America and China stand to see the biggest economic gains with AI enhancing GDP by 26.1% and 14.5% in 2030, respectively, equivalent to a total of US$10.7 trillion and accounting for almost 70% of the global economic impact. Beyond North America and China, other countries such as those across Europe and the more developed countries in Asia are also likely to experience significant GDP gains of around 9.5%–11.5% of GDP by 2030 (Rao et. al., 2017). Although the adoption of AI is projected to be slower in these countries than in the North American region, the potential for automation is high in Europe, while the marginal impact of AI technologies on productivity is particularly high in developed Asian economies, as is projected investment in workforce-augmenting technologies.

On the consumption side, fast adopters of AI are likely to see the greatest gains (North America and Northern Europe), although China will see a disproportionately large share of the benefits due to the slightly lower level of competition in its landscape, which increases the marginal impact of firm entry on prices, discussed in more detail below. Latin America and other less developed markets are expected to lag behind somewhat, though despite lower uptake of AI they are still expected to see GDP gains of approximately 5% in 2030 (see figure 2).

THE RISKS OF AI

For all the enormous potential AI offers for building a sustainable planet for future generations, it also poses short and long-term risks. These can be divided, broadly speaking, into six categories with varying impacts on individuals, organizations, and society (table 1).

Performance Risks

Like any other software system, AI

FIGURE 1:
GLOBAL GDP IMPACT OF AI THROUGH 2030

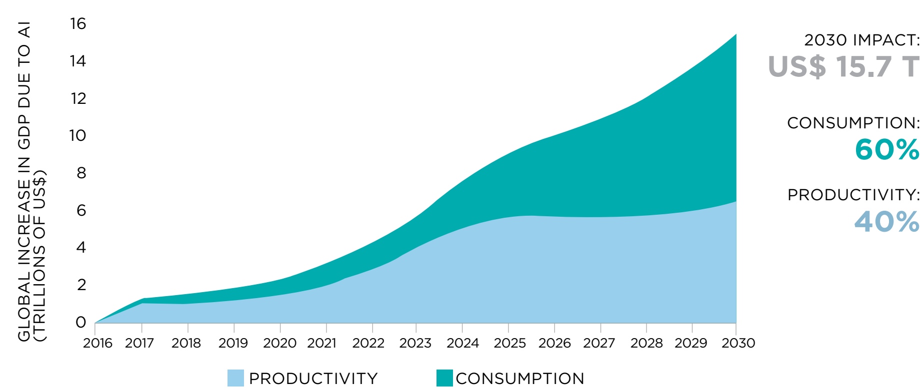

Source: PwC Data & Analytics

5%+

THE GROWTH IN LAC'S GDP DUE TO AI BY 2030

systems need to be verified and validated using standard methodologies. However, AI systems—particularly ML systems—differ significantly from standard software systems. Broadly there are two phases to building a machine ML system (Dietterich, 1988; Hall, Phan, and Ambati, 2017). First, the developer trains the system by providing large volumes of input and output data. For example, if you want to have the ML system to identify a cat from a number of images, the developer feeds hundreds and thousands of images, a subset which has cats in them that are clearly marked. The ML system then learns to identify features that are unique to cats. Once the system has been adequately trained, it is deployed in production mode, where the ML system will identify if there is a cat in any new image it is presented with.

The fundamental difference between traditional systems is that there is no line-by-line code for someone to verify. Instead, we need to ensure that the data provided is representative, that there is no bias in the data, and that we understand how the system is identifying the features and how it is making the recommendations. The difficulty of doing these on many of the ML algorithms makes them a black box, making it difficult to ascertain whether the performance or outputs of AI algorithms are accurate or desirable. The emerging field of explainable AI (XAI) research (Gunning, 2016) aims to create new AI methods that are accountable to human reasoning. But this field is still in its early days. Meanwhile, ongoing research aims to reduce "model bias" resulting from biases in training data and to increase the stability of model performance. As AI solutions are deployed, one unintended consequence is the over-reliance on AI algorithms with variable performance (Goodman and Flaxman, 2016). It is essential that humans stay in the loop on auditing algorithm outputs to mitigate these unintended biases and wider performance risks.

For example, a number of banks and insurance companies are using ML models to make decisions around granting consumers loans, credit cards, or insurance policies. If the data used by these organizations is biased or is not representative of the entire population or the ML system is unable to explain the logic of its recommendation in a manner that consumers can understand, then in the widespread adoption of these techniques, we run the risk of consumers losing trust in the system's recommendations.

Security Risks

Misuse of AI via hacking is a serious risk, as many algorithms being developed with good intentions (for example, for autonomous vehicles) could be repurposed for harm (for example, for autonomous weaponry). This raises new risks to global safety (Brundage et. al., 2018).

Good governance is required to build explainability, transparency, and validity into the algorithms (Easterbrook, 2010), including drawing lines between beneficial and harmful AI (Holdren and Smith, 2016). ML models (especially deep learning) can also be duped by malicious inputs known as "adversarial attacks." For example, it is possible to find input data combinations that can trigger perverse outputs from ML models, in effect hacking them.

For example, hackers could access automated warning systems, distributed energy grids, or connected autonomous transportation platforms and cause regional disruptions. Appropriate governance will be required to ensure human and environmentally friendly AI and prevent misuse. Misuse of AI could also occur when systems fall into the wrong hands. For example, poachers could profit from AI-enabled endangered-animal tracking tools meant for conservation efforts.

Control Risks

Some AI systems work autonomously and interact with one another, creating machine-centered feedback mechanisms that can cause unexpected outcomes. Semi-autonomous and autonomous vehicles, sensor-enabled heavy machinery, drones, robots and a number of other devices and equipment will increasingly have AI embedded within them. Human inability to take control of these semi-autonomous or autonomous systems introduces major control risks (Brundage et. al., 2018).

In 2010 a financial crash was caused by the interactions of multiple AI bots speed-trading, which created artificial market inflation. Similarly, hackers have demonstrated how they can take control of vehicles and remotely control them. In the wrong hands, these could lead to significant risks to people and property. Proactive control, monitoring, and safeguards are necessary to catch these issues before they become a problem. Control of AI systems or emergency human intervention needs to be factored into the design of these systems.

Economic Risks

As companies adopt AI, it may alter the competitive landscape, creating winners and losers. Those able to improve their decision-making most quickly through AI may find the bene-

FIGURE 2
GROWTH BY GEOGRAPHIC AREA

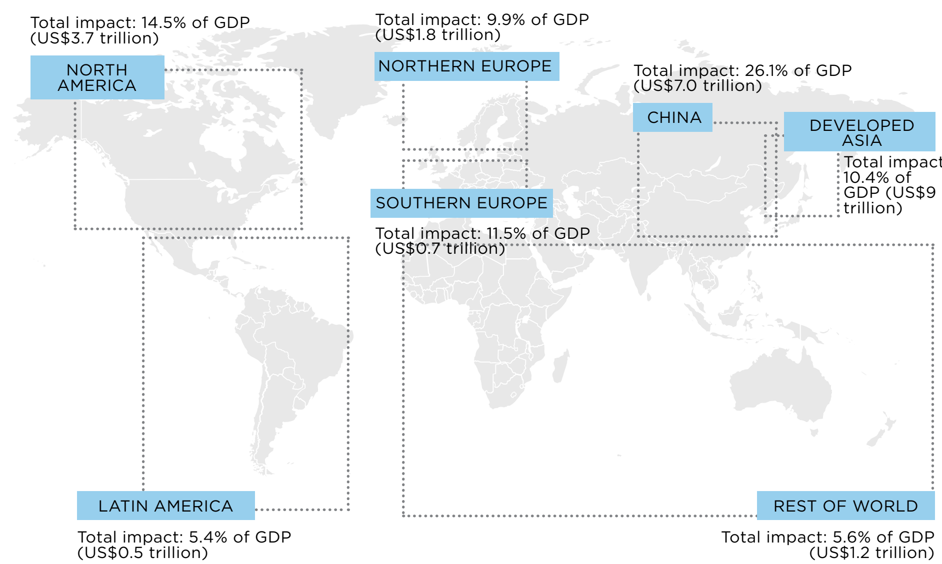

Source: PwC Data & Analytics

fits accelerate very quickly, while slower adopters may be left behind. Companies that struggle in the AI transition may be forced to reduce investment, possibly impairing their profitability and eventually their existence. Given the accelerating returns on cognitive capital (the combination of human and machine intelligence), the first movers with the right data and experts can quickly monopolize their market. Given the global nature of the digital world, this could very quickly result in a race for global supremacy, forcing governments to intervene to protect their local industries and potentially paving the way for more protectionism and less globalization.

For example, increased productivity from automation, plus rising consumption from improved personalization, product design, and AI-informed marketing will change the number of people required to deliver these goods and services and could also skew the nature of the skills required to survive in the emerging AI world.

Social Risks

Large-scale automation threatens to reduce employment in transportation, manufacturing, agriculture, and services, among other sectors. Higher unemployment rates could lead to greater inequality in society. In addition, algorithms designed by a subset of the population at a national and global level have the potential for unconscious bias, possibly leading to results that marginalize minorities or other groups. Autonomous weapons also pose a significant threat to society, possibly permitting bigger, faster conflicts. Once unleashed, this might lead to rapid and significant environmental damage, even to a doomsday scenario where weaponized AI presents an existential risk to humanity.

For example, autonomous trucks and cars, along with energy-efficient Internet of Things manufacturing, offer considerable environmental benefits but could also lead to a considerable loss of employment (Goldman Sachs estimates that the US alone will lose an estimated 300,000 jobs per year when AI saturation peaks). Regional economic decline and widening social inequality and unrest could also follow.

Ethical Risks

The ethical and responsible use of AI involves three main elements: the use of big data; the growing reliance on algorithms to perform tasks, shape choices, and make decisions; and the gradual reduction of human involvement in many processes. Together, these raise issues related to fairness, responsibility, equality, and respect for human rights. Additionally, while biased AI outcomes can raise significant privacy concerns, many insights and decisions about individuals are based on inferred group or community attributes. Accordingly, consideration of the harm AI could do must be framed beyond the individual level and recognize that privacy is not the only issue.

For example, an autonomous vehicle in the future is likely to face situations where it may need to make moral decisions (if faced with two choices—potentially killing one passenger in the vehicle or two on the road, what decision will it make?). Humans make different choices based on their values and how to impart those values to machines or at least align the values of machines with those of humans is a challenging problem.

CREATIVE DESTRUCTION

AI has the potential to fundamentally disrupt markets through the creation of innovative new services and entirely new business models. We've already seen the creative destruction of the first wave of digitization.

With the next wave of AI, some of the market leaders in ten or even five years' time may be companies you've never heard of. In turn, some of today's biggest commercial names could be struggling to sustain relevance or have even disappeared altogether, if their response has been too little or too late.

Tomorrow's market leaders are likely to be exploring the possibilities and setting their strategies today. Companies should be actively developing their AI strategy and enhancing their AI capabilities (Rao, 2017). We recommend that companies focus on the following areas:

Develop an AI strategy aligned with their overall business strategy: They should integrate AI into their existing digital and analytics plans; decide which businesses to disrupt and which to enhance; consider new business models based on improved productivity; and plan long-term investments in autonomous intelligence.

Develop enterprise-wide AI capability: redesign products and services to incorporate ML; use AI to upgrade your most critical distinctive capabilities; use automation to improve current decisions; automate existing business processes or develop new ones; and recruit engineers and other professionals who understand AI.

Institutionalize their portfolio of AI capabilities: embed AI throughout business processes; embrace cloud platforms and specialized hardware; and foster a decision-making culture open to AI-supported ideas.

Ensure appropriate governance: establish clear policies with respect to data privacy, decision rights, and transparency; set up governance structures to monitor possible errors and mitigate the risks from AI; develop explainable AI and set up governance processes to monitor the efficacy of these explanations; consider the impact on employment; and invest in developing the workforce that AI will complement.

As we anticipate the tremendous changes that AI will bring, it's important to remember that we are only at the beginning of the AI revolution. We expect further disruption and transformation in all economies, geographies, and sectors. While AI's economic impacts might not be fully realized for a decade or more, businesses must begin making the right strategic moves today.

REFERENCES

Brundage, M., Avin, S., Clark, J., et al. 2018. "The Malicious Use of Artificial Intelligence: Forecasting. Prevention, and Mitigation." University of Oxford. arXiv:1802.07228.

DARPA. 2016. "Broad Agency Announcement: Explainable Artificial Intelligence (XAI)." DARPA-BAA-16-53. Arlington: Defense Advanced Research Projects Agency (DARPA), Information Innovation Office.

Dietterich, T.G. 1988. "Artificial Intelligence." Minsky 1988 (2017): 3–24.

Easterbrook, S. 2010. "The Difference between Verification and Validation." *Serendipity*. November 30.

Gillham, J., Rimmington, L., Dance, H., et al. 2018. "The Macroeconomic Impact of Artificial Intelligence." PwC White Paper.

Goodman, B. and Flaxman, S. 2016. "EU Regulations on Algorithmic Decision-Making and a 'Right to Explanation'." *AI Magazine* 38 (3): 50–57.

Hall, P., Phan, W., and Ambati, S. 2017. "Ideas on Interpreting Machine Learning." O'Reilly Media. March 15.

Holdren, J., and Smith. M. 2016. "Preparing for the Future of Artificial Intelligence." Washington, DC: Executive Office of the President of the United States. National Science and Technology Council Committee on Technology. US Federal Government.

Kelnar, D. 2016. "The Fourth Industrial Revolution: A Primer on Artificial Intelligence (AI)." Medium, MMC Writes. December 2.

Rao, A. 2017. "A Strategist's Guide to Artificial Intelligence." Strategy + Business 87: 2–10.

Rao, A., Verweij, G., and Cameron, E. 2017. "Sizing the Prize: What Is the Real Value of AI for Your Business and How Can You Capitalise?" PwC White Paper. PwC.

Schwab, K. 2016. The Fourth Industrial Revolution. Geneva: WEF.

Peter Diamond
Nobel Prize in Economics

“Technology helps countries to build flexible rules”

PETER DIAMOND, WINNER OF THE NOBEL PRIZE IN ECONOMIC SCIENCE, IS ONE OF THE LEADING GLOBAL EXPERTS ON SOCIAL SECURITY, LABOR MARKETS, AND THE ANALYSIS OF MARKET FAILURE, ECONOMIC PROBLEMS THAT REQUIRE GOVERNMENT INVOLVEMENT TO BE SOLVED. IN THIS INTERVIEW WITH INTEGRATION & TRADE, DIAMOND EXPLAINS THAT THE SPEED OF TECHNOLOGICAL CHANGE IS FORCING GOVERNMENTS TO INCREASE THEIR RESPONSE CAPACITY TO ENSURE THEY ARE NOT LEFT BEHIND BY THE PRIVATE SYSTEM AND NEW CHALLENGES. HE ALSO PRAISES THE CHILEAN PENSION SYSTEM AND ARGUES THAT THE RULES ON TRADE AND GOODS FOLLOW A DIFFERENT LOGIC TO THE RULES AROUND CAPITAL MARKETS.

How will new technologies impact employment?

I like to think about it in a long-term context. Obviously, we have had technological knowledge advances that have had a major impact on the economy for several centuries now. The first lesson that comes from all these is that it is very important for the government to be responsive to the issues raised by new technologies. We are familiar with the role of government in regulating business. The way businesses pursue profits may call for consumer protection, worker protection, or environmental protection. We know the nature of risks in the economy and the role of the government in dealing with risks that the market doesn't deal well with. It is inevitable is that business will move much faster in response to opportunities than governments will in response to issues they could make better.

> LONG-TERM SUSTAINABILITY DEPENDS ON POLITICS, NOT ECONOMICS

How can we improve the way labor markets function?

That is always an issue. At the time of the financial crisis, Paul Samuelson talked about how financial engineering was the cause of the financial crisis. And Bob Merton said that it's always the case that technology runs ahead of the safety regulations that go with it. That was very much the case with derivatives. So it's nothing new. If through AI we have drones that will be fighting against each other without people, among many other things, then the question is how this will impact both governments and businesses. Things seem to be happening much faster, and the problem is that we have been neglecting these changes over the last 30 years. We haven't invested in education or in infrastructure and we have been doing perverse things about income distribution issues. Now, we need to focus on where we can get the revenue we need for income distribution. The emergence of these technologies doesn't mean it's necessarily negative, but the key question we need to address is how we can get a lot more value for people.

Do you think unemployment will increase as a result of automation?

If we focus on the labor market,

first, it would be wrong to focus on unemployment as the issue. We have had this situation before, and yes, you get a short-run impact on employment. However, labor markets tend to respond and adjust (when they are allowed to). The issue is much more about the extent to which the education system prepares people for the opportunities that will emerge. What we are seeing here is a change in relationships, a change in relative power, a change in organization, which is cause for rethinking regulations and the rights that workers have, who's an employee and who's a contractor, what are the obligations of a contractor. These problems force us to rethink many things.

WE NEED MAJOR COORDINATION ACROSS COUNTRIES ON TAXATION ON MULTINATIONAL CORPORATIONS

How can we organize social security systems to take into account longer life expectancy and freelancing?

The answer to that varies from country to country. There are some countries that have set up systems that will stay in place and will adapt to changes in life expectancy, changes in income distribution, and other factors like that. Chile has designed a good system and since they have a well-functioning government, they have the ability to make it work, being able to adapt regulations frequently, adjusting them to specific necessities. Other countries have systems that are much more expensive and they haven't changed anything about them for years and years. All of Latin America has the problem that the formal sector is half the labor force at its best, so the pension system, which is built around formal employment, is not dealing with the whole population. The solidarity pillar in Chile is a way to address that gap. And there are two aspects to sustainability: first, what level makes sense in relation to what the economy can afford. And second, which is a more political issue, is how much revenue will be collected. Sustainability is therefore really a political question, not an economic question. In Sweden, for example, they make automatic adjustments. If the pension system is projected to be unsustainable, benefit cuts start right away. You still have the issue of how large benefits will be, but they don't have the shock of a sudden movement. The idea is that sustainability is something you can build with automatic rules, and the hope is that you will do a better job in designing automatic rules than would be the outcome with just periodic responses to crisis.

Do you think that artificial intelligence could help design better standards and public policies?

Yes, we will have the ability to run more complex pension systems thanks to computerization and communications. Through these technologies, you can spread risk in a more intelligent way. But the other dimension is that there would be a large section of the population in those countries, probably a majority, who are not prepared to make those risky decisions. As such, the whole framework needs to be designed well in advance. One of the things that has received a lot of attention is the premium pension system in Sweden. People have recognized that automatically being assigned to a system is a way of getting people to take part even if they don't want to spend time deciding what stocks to buy. Consequently, nowadays 99% of people who are entering the system, mostly young people, are in the default option. This is why system design is very important. You can adjust the system in a more sophisticated fashion using technology to tailor it to the risks that people are willing to take on. You can have multiple pieces, as in Chile and as is also common in advanced countries, which take a nonlinear approach to income distribution. In most countries, the system has several plans, one of which is linear and the other is flat. And then there may be some links between the two, as in Chile, so there's an offset. In the Netherlands, there is a flat benefit for people over 65 years of age. It's set above the poverty line, so elderly poverty is almost nonexistent. Obviously, it takes tax revenue to do that, and they have a dedicated piece of income tax that is meant to cover that. Again, you need to have a tax in place that is linked to a kind of spending that people care about.

IN 1982, PETER DIAMOND, THEN A YOUNG PROFESSOR AT MIT, PUBLISHED AN ARTICLE THAT WOULD SHAKE UP THE ACADEMIC WORLD AND WOULD COME TO BE KNOWN AS THE "DIAMOND COCONUT MODEL." THE TEXT PROVIDED RIGOROUS PROOF FOR WHY PEOPLE WOULD CLIMB UP PALM TREES TO LOOK FOR COCONUTS ONLY IF THEY BELIEVED OTHER PEOPLE WOULD ALSO DO SO, THUS LAYING THE FOUNDATIONS FOR ANALYZING COORDINATION PROBLEMS. THIS WAS ONLY THE BEGINNING OF A PROLIFIC ACADEMIC CAREER: DIAMOND HAS PUBLISHED OVER 12 BOOKS AND 150 ARTICLES. THE FOLLOWING IS A VERY BRIEF SELECTION:

"AGGREGATE DEMAND MANAGEMENT IN SEARCH EQUILIBRIUM." JOURNAL OF POLITICAL ECONOMY 90 (5): 881-894. 1982.

"PAIRWISE CREDIT IN SEARCH EQUILIBRIUM." QUARTERLY JOURNAL OF ECONOMICS 105 (2): 285-319. 1990.

"THE FLOW APPROACH TO LABOR MARKETS." AMERICAN ECONOMIC REVIEW 82 (2): 354-359. WITH O. BLANCHARD. 1992.

What type of coordinated action can countries take?

One word that we haven't mentioned yet is "globalization" and the response to technology in different international environments. For example, on the issue around tax revenue, we need major coordination across countries on taxation on multinational corporations. We also need to shut down tax havens, something which is politically difficult to do. Another big issue is foreign investment and the rules that go with it so you don't get the kind of boom and bust that is so common in Latin America. The first thing to think about is that the rules on trade and goods follow a different logic than the rules around capital markets. The idea that the capital market should be equally open both ways is a bad idea. Countries do need foreign direct investment, but they need it in an environment where there is protection for workers, consumers, pensions, and the environment. To be honest, I don't know how to get there, I'm not a political scientist. What we need in the regulatory realm is more good regulation and less bad regulation. In other words, more or less is not the question, but the quality of it.

How to *Facilitate* Trade Negotiations

Welber Barral
Barral M Jorge and Associates

Gabriel Petrus
International Chamber of Commerce in Brazil

ARTIFICIAL INTELLIGENCE COULD BE USED TO REVOLUTIONIZE THE CURRENT RULES OF PLAY FOR INTERNATIONAL RELATIONS, AS THEY CAN REDUCE THE COMPLEXITY OF TRADE NEGOTIATIONS AND INCREASE TRANSPARENCY AND ACCESS TO INFORMATION.

On the eve the Fourth Industrial Revolution, the extent to which disruptive technologies are reshaping the way our national and global institutions are organized, and the speed at which they are doing so, are becoming increasingly evident (Schwab, 2016)[1].

It is worth noting the exponential speed at which advanced manufacturing has been increasing industrial productivity over the last few years. Robotics and automation are not only bringing more efficiency to the material world but, at the end of the day, they are also influencing and changing the way key areas of our society and governments make sensitive decisions.

Augmented intelligence tools are built over deep learning and cognitive systems—known by the general public as artificial intelligence—and they are already in widespread use by different sectors and industries.

In an op-ed article for the Intelligent Tech & Trade Initiative—a research[2] project launched by ICC Brazil aiming to identify applications for artificial intelligence tools in trade negotiations and transactions—IBM CEO Ginni Rometty recognizes the tremendous potential of AI systems in our society. These include "predicting risk in financial markets, anticipating consumer behavior, ensuring public safety, managing traffic, optimizing global supply chains, personalizing medicine, treating chronic diseases, and preventing pandemics" (Rometty, 2017: 15).

Moreover, there are unexplored frontiers for the use of AI to benefit global advancements, including international trade: open markets are key to economic growth and shared prosperity. AI is not intended to replace humans or human will, which essentially distinguishes the human condition from that of other living beings (Arendt, 1958).

One might use different arguments now, but the same distinction still applies to human beings when compared to robots—something which Hannah Arendt surely could not have predicted in the 1950s.

Augmented intelligence has huge potential to replace ill-oriented or inefficient decisions by shedding light on complex decision-making processes and making them more efficient and cost-effective. In this sense, we see a promising opportunity for governments and multilateral organizations to apply AI to international trade negotiations, transactions, and operations.

As well as making trade more efficient, these tools have the potential to play a game-changing role in the current state-of-the-art of international affairs. As the complexity of trade negotiations will be diminished and transparency will increase, it will more difficult for protectionist policies or antiglobalization rhetoric to provide shelter for misleading economic arguments.

This article will explore the potential application tracks of AI in trade negotiations.

COMPLEX NEGOTIATIONS

Reaching successful trade agreements are among the most complex and lengthiest endeavors for policy-makers worldwide in both developed and developing countries.

Regional and bilateral negotiations usually take longer than the terms of office of the governments that are in charge of them. The MERCOSUR—European Union negotiations toward a bilateral agreement, for instance, have been going on for more than 15 years.

Even for countries that have historically played the role of trade liberalization champions, free trade agreement negotiations have rarely taken less than three years to conclude, from launch to signing.

According to a study conducted by Caroline Freund and Christine McDaniel (2016), "on average it takes one and a half years to negotiate an FTA with the United States but over three and a half years to reach the implementation stage." Taking the US as a case study, the research reveals that trade negotiations with countries with higher trade shares are the lengthiest to conclude.

The logic underpinning these findings is that a higher trade share means that more sectors and more products are involved in the process, which means, in other words, that more disaggregated data is involved. As pointed out by ITTI, this information includes the taxes levied on products under negotiation, the list of products, the rules of origin framework, the potential for supply chain integration, the volume of exports and imports between the parties involved, previous trade agreements, other statistics, and more.

Moreover, in other research carried out by Christoph Moser and Andrew Rose (2012), the durations of regional and multilateral trade negotiations were empirically modeled. When it comes to regional trade negotiations, an analysis of a historical series of more than 30 trade agreements revealed that "[these negotiations] are more protracted when there are more countries at the negotiation table."

Negotiations tend to be considerably more complex at the multilateral levels. Moser and Rose (2012) established an interesting comparison based on the duration of GATT/WTO trade liberalization rounds. According to these authors, "the length of time between the start of negotiations and their completion has grown consistently with the number of participants."

The study points out that "the 23 participants in the first (Geneva) round of GATT negotiations took only six months to conclude a deal that reduced 45,000 tariffs." Nowadays, with over 150 members of the WTO, the data analyzed points to the conclusion that the substantial number of involved parties makes negotiations considerably more difficult. The number of stakeholders involved is thus another key factor in the complexity of these trade negotiations.

TABLE 1
DURATION OF GATT/WTO ROUNDS

ROUND	INITIATED	COMPLETED	PARTICIPANTS	DURATION
GENEVA	APRIL 1947	OCT. 1947	23	6 MONTHS
ANNECY	APRIL 1949	AUGUST 1949	13	4 MONTHS
TORQUAY	SEPT. 1950	APRIL 1951	38	7 MONTHS
GENEVA II	JANUARY 1955	MAY 1956	26	16 MONTHS
DILLON	SEPT. 1960	JULY 1962	26	22 MONTHS
KENNEDY	MAY 1964	JUNE 1967	62	37 MONTHS
TOKYO	SEPT. 1973	NOV. 1979	102	74 MONTHS
URUGUAY	SEPT. 1986	APRIL 1994	123	91 MONTHS
DOHA	NOV. 2001		153	123 MONTHS

Source: GATT (1980) and WTO (2018).

Another factor behind this complexity that Moser and Rose (2012) highlighted as causing overly long trade negotiations is the economic profile of the countries involved. Negotiations between more open, richer countries finish more quickly than negotiations between more closed, poorer economies. Consequently, we can assume that asymmetry of information is another key variable in the protraction of trade negotiations.

What these researchers have brought to light is that data-driven variables influence the length of trade negotiations much more than politically driven variables do. As a consequence, augmented intelligence technologies have the potential to tackle the roadblocks that make trade negotiations less efficient than they might be.

The ITTI has highlighted another important consequence of the application of AI to trade negotiations: "access to these technologies for both developed and least developed countries will be a turning point in how nations create trade policies and could minimize the influence of interest groups and politics when making trade decisions, as well as boost international trade as a whole" (ITTI, 2017).

To raise awareness of the potential importance of adopting these technologies, ambassador Álvaro Cedeño summed up the situation: "the world has changed dramatically, mainly due to technological progress. The acceleration of microprocessing and memory storage capacity and internet connectivity speeds has been particularly dramatic in the last ten years since the rise of the smartphone. Yet trade policy is still being negotiated in a rudimentary and widely ineffective manner" (Cedeño Molinari, 2017).

45,000
TARIFFS WERE CUT BY THE 23 GATT MEMBERS IN JUST 6 MONTHS

TRADE FACILITATION

Based on predictive analysis using machine learning models, augmented intelligence offers an efficient instrument for responding to the complexity variables identified above and promoting more efficiency in trade negotiations.

Besides, more efficient data processing will become more accessible for all parties involved and the focus of trade negotiators' key decisions will be more technical, rather than politically oriented. As ITTI (2017) has observed, "augmented intelligence is adding a human touch to big data. Rationality is assisted and leaves more space for decisions based on values."

In this sense, Álvaro Cedeño highlights the importance of adopting digital technologies that accelerate the application of international trade and trade negotiation simulations: "This could render results that would serve as evidence and metrics at the service of decision- and policy-makers (...) simulations could eventually lead the institution to the adoption of artificial intelligence as a resource for employing self-generated data which would help visualize the causes and consequences of the extraction of raw materials, production, manufacturing, shipping, exports and imports, consumption, and disposal of goods" (Cedeño Molinari, 2017).

Another component of trade negotiations that can be transformed by AI are the so-called small technical issues that actually have the potential to hinder or block these processes. A simple example is the determination of rules of origin, an arcane topic that has regularly delayed many trade negotiations.

According to Eric Siegel (2016), for cognitive computing, no detail is too small to make it not worth considering. As a result, if augmented intelligence is used in a trade negotiation, no microrisks (or micro-information, which might be important for a certain economic sector interested in a specific chapter of a trade negotiation) will go unnoticed. This will be extremely helpful for making trade negotiations an easier endeavor for decision-makers.

Another benefit that comes from the use of AI applications is the capacity to process the huge amount of data on existing trade agreements and their impact on ongoing negotiations. One of the ITTI's findings is that "automated data analysis allows us to track several years of decisions that can help us understand the roots of other parties' arguments. Negotiation can be faster as objective issues can be solved with more documented information" (ITTI, 2017: 49).

Finally, while mechanization has meant that human beings have to work less (Chase, 2016), AI will mean that trade negotiators will have more time to focus on key political aspects of trade negotiations rather than on the technicalities that often act as roadblocks to reaching effective agreements.

In other words, they will have more extra time to tackle issues like the political relevance of an agreement for domestic politics, or the transitioning required for a sensitive industrial sector, or the strategic impact for a long-term bilateral relationship, all issues that require political analysis rather than technical evaluation.

INCREASING COMPETITIVE ADVANTAGES

Another potential use of AI tools for countries during trade negotiations is to allow each of them to explore and expand their competitive advantages. According to ITTI (2017: 43), "using augmented intelligence to make trade decisions could allow countries to expand their competitive advantages by helping them better inform their trade policies."

The ITTI corroborates the previous arguments presented in this article that it is time-consuming for government officials and trade negotiators to gather all data necessary to start a trade negotiation and see it through. Thus, it would not only reduce the complexity of trade negotiations but would also speed up the process of gathering data and help trade negotiators to define the best strategy for successfully concluding a negotiation.

ITTI also asserts that AI will be useful for analyzing potential barriers to trade and identifying which tools would fit the purposes of both parties and suggest win-win situations between two poles in a trade negotiation that would lead to the successful conclusion of this.

ADOPTING AI TOOLS

As is the case with every disruptive initiative, there are a few challenges that the public and private sectors need to tackle if they are to adopt AI tools for use in trade negotiations. We believe these are not technological challenges—AI has been used to transform much more sensitive fields, such as healthcare and public security—but rather governance challenges.

They can be divided into two areas: (1) setting global rules of play and (2) establishing a global trade database that will feed these platforms (and their algorithms) with reliable information.

There are currently multiple institutions managing their own trade data repositories without interacting with each other: UN Trade, the World Bank, the WTO, UNCTAD, and the International Trade Center, to name just a few. A cloud-based, integrated database thus needs to be created to ensure that the same level of information can be accessed by the different AI tools built and used by trade negotiators.

The ITTI (2017: 46) proposes that "an international cloud-based resource with information on all international agreements ratified worldwide could inform negotiators on how to best proceed within new trade negotiations."

According to Álvaro Cedeño, "this would require a great deal of trust among trading partners, a great deal of transparency in order to allow the most up-to-date information to be the basis for those decisions, and domestic and international regulatory frameworks that are seamlessly synchronized to enable such frictionless trade" (ITTI, 2017: 62).

Another challenge is to establish common parameters for what is considered a "successful" or "fair" outcome of a particular trade negotiation. As the ITTI (2017: 50) states, "achieving uniformity and consensus on what parameters to adopt will not be easy, given nations' different political economy interpretations of what 'good' trade negotiation results actually mean."

In parallel, it must be ensured that all countries—including the least developed countries (LDCs)—have access to this technology in order to prevent the trade and tech divide. As the ITTI (2017: 44) argues, "for LDCs, the biggest roadblock will be access to these innovative technologies, though investing in them could be greatly beneficial." The good news is that cloud-based solutions are becoming increasingly more accessible than ordinary hard infrastructure, which tends to be costlier for LDCs.

To be effective, these tools must be used by a significant number of countries that are engaged in trade negotiations. Widespread use is a key aspect of their success. The more trade data the tool is given the more comparisons the platform can make and the more effective it will be for future negotiations. This is a core aspect of what is defined as deep learning processes.

THE ROAD AHEAD

The ITTI is a global research project led by ICC Brazil to bring together technology and business leaders, negotiators, and scholars to debate and devise ways for blockchain (the trust ledger) and augmented intelligence to positively impact global trade.

The ITTI's ultimate objective is to nourish debate between the technology community, trade negotiators, business leaders, and scholars on how to better pursue a constructive trade agenda. Mindful of both national and multilateral specificities, the ITTI aims to counter the deglobalization forces now affecting international trade.

The ITTI was officially launched by ICC Brazil on October 2017 during the WTO Public Forum in Geneva and it released its first discussion paper in December 2017 at the WTO Ministerial Conference in Buenos Aires. Its next steps include the development of a global app track for trade negotiators based on artificial intelligence platforms.

Therefore, global rule-making institutions such as the WTO—and the ICC—have a role to play in raising awareness of these mechanisms, harnessing their potential to promote more trade agreements at the bilateral, regional, and plurilateral levels.

LESS SUBJECTIVITY

The use of AI is already revolutionizing policymaking (Wigglesworth, 2018), and will enable nations to be led by smart governments. This is already a becoming reality for certain key areas of government that are using these platforms to leverage processing capacity for large amounts of economic and social data, based on much more accurate models to enable economic and policy decisions.

This article has explored the tremendous potential of these technologies to empower the decision-making process across borders and between different national and international stakeholders. Their impact on international trade has great potential to produce smart trade negotiations.

As mentioned before, AI will not replace political decisions, but it will make them more transparent, increasing the capacity of trade negotiations to process large amounts of complex data efficiently.

Consequently, AI will decrease the level of subjectivity in trade negotiations, neutralizing ill-oriented economic arguments that have been underpinning the protectionist rhetoric and populist stances.

AI may considerably ease negotiations around trade agreements, helping to realize all of trade's potential to contribute to economic growth. At a historical moment in which the idea that trade is a key element for economic growth has diminished, the use of disruptive technologies may be also a tool to prevent the undermining of the multilateral trading system.

NOTES

[1] Welber Barral is a former secretary of foreign trade of Brazil (2007–2011) and is a managing partner of Barral M Jorge and Associates, an international trade consulting firm. Gabriel Petrus is the executive director of the International Chamber of Commerce in Brazil (ICC Brazil).

[2] For more on the Intelligent Tech and Trade Initiative, see the "Road Ahead" box in this article.

REFERENCES

Arendt, H. 1958. *The Human Condition.* Chicago: University of Chicago Press.

Cedeño Molinari, Á. 2017. "Smart Trade Multilateralism." In: *Intelligent Tech and Trade Initiative. Building ITTI: A Discussion Paper*. São Paulo: Câmara de Comercio Internacional de Brasil.

Chace, C. 2016. *The Economic Singularity: Artificial Intelligence and the Death of Capitalism.* Three Cs. Kindle Edition.

Freund, C., and McDaniel, C. 2016. "How Long Does It Take to Conclude a Trade Agreement With the US?" Peterson Institute for International Economics, Trade & Investment Policy Watch. Washington, DC.

GATT. 1980. "Loose-Leaf System for the Schedules of Tariff Concessions: Abbreviations for GATT Legal Instruments." Committee on Tariff Concessions. TAR/W/6. May 27.

ITTI. 2017. *Intelligent Tech and Trade Initiative. Building ITTI: A Discussion Paper*. São Paulo: Câmara de Comercio Internacional de Brasil.

Moser, C., and Rose, A. K. 2012. "Why Do Trade Negotiations Take So Long?" *Journal of Economic Integration* 27 (2): 280–290.

Rometty, G. 2017. "General Landscape of Technology." In: *Intelligent Tech and Trade Initiative. Building ITTI: A Discussion Paper.* São Paulo: Câmara de Comercio Internacional de Brasil.

Schwab, K. 2016. *The Fourth Industrial Revolution.* Geneva: World Economic Forum.

Siegel, E. 2016. *Predictive Analytics: The Power to Predict Who Will Click, Buy, Lie, or Die.* Hoboken: Wiley.

Wigglesworth, R. 2018. "Can Big Data Revolutionise Policymaking by Governments?" *Financial Times.* January 31.

WTO. 2018. "The GATT years: from Havana to Marrakesh." Available at: https://www.wto.org/english/thewto_e/whatis_e/tif_e/fact4_e.htm

Dave Donaldson
Massachusetts Institute of Technology
(MIT)

“The impact of **AI on the economy** will depend on each country’s **openness to trade**”

Dave Donaldson is one of a select list of economists. This young professor at the Massachusetts Institute of Technology (MIT) was awarded the American Economic Association’s John Bates Clark Medal, which is given to the best US economist under the age of 40. Paul Samuelson, Milton Friedman, Robert Solow, James Tobin, Kenneth Arrow, Gary Becker, Joseph Stiglitz, and Paul Krugman, among other Nobel laureates, also received the award early in their careers. Donaldson is an expert on international trade and was one of the speakers at the Economics of Artificial Intelligence Conference, which brought major international experts together in Toronto in 2017 to discuss the future of the economics of algorithms. In this interview, he analyzes the impact of artificial intelligence (AI) on trade policy.

What conclusions did the economists who attended the conference in Canada reach?

It was a hugely eye-opening introduction to all the breakthroughs that AI is already making in different areas of life and will, I’m sure, continue to make. The likelihood that AI will impact economic life seems very high—much too high to ignore.

What implications does AI hold for trade policy?

Countries could use trade policy either as a particularly targeted form of domestic intervention or as an attempt to change their terms of trade in a way that favors themselves and by nature harms other countries. But both of these are unlikely to be good policy decisions. Domestic interventions are far better dealt with via purely domestic policy—which country uses tariffs to try to get polluters to pollute less? Most countries are unlikely to be able to do much to change world prices and hence their terms of trade, so when something in the economy changes, like AI getting developed, my first reaction is always the same: “trade policy probably shouldn’t adapt.” Again, if you’re thinking of fixing a domestic problem (such as concerns about rising inequality or lagging regions), then those can be fixed far better via domestic policies. And if you’re in a relatively small country and you’re thinking that you can use trade policy to offset the effects of AI on world prices, then you’re probably kidding yourself.

How important is geography if knowledge spillovers are to take place?

The truth is that we don’t really know. Nobody knows. There is some evidence that is consistent with knowledge spilling over. But that evidence is by no means conclusive. We all know that knowledge and information don’t just flow everywhere all the time. If that were true, then people wouldn’t be able to process it all. It has to do with the selectivity of the process by which knowledge leaves its source and is absorbed by the destination. Obviously, it’s plausible that the flow might become costlier at a distance, but it’s also true that today we academics write papers and we put them on the internet, and that flow of knowledge is no more local than global. To take an example from academic research outside of economics:

After receiving a master's in physics from the University of Oxford, Dave Donaldson decided to switch his focus to economics. He received his Ph.D. from the London School of Economics and began a prolific publishing career that has included articles in all the major economic journals. His most noteworthy publications include:

· "Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade." *American Economic Review* 107 (3): 633–689. With R. Adao and A. Costinot. 2017.

· "The View from Above: Applications of Satellite Data in Economics." *Journal of Economic Perspectives* 30 (4): 171–198. 2016.

· "The Gains from Market Integration." *Annual Review of Economics* 7: 619–647. 2015.

· "Ricardo's Theory of Comparative Advantage: Old Idea, New Evidence." *American Economic Review*: Papers and Proceedings 102 (3): 453–458. With A. Costinot. 2012.

· "What Goods Do Countries Trade? A Quantitative Exploration of Ricardo's Ideas." *Review of Economic Studies* 79 (2): 581–608. With A. Costinot and I. Komunjer. 2012.

Google wanted to set up centers of expertise in various aspects of AI and they wanted certain famous, high-profile, high-impact researchers to head these centers. So, the question arose as to where to locate these offices, and apparently these high-profile researchers were reluctant to move, so Google just built its centers around them. The company now has one in London, England, and another in Edmonton, Canada—two locations that are not exactly close to the main Google headquarters in Mountain View, California. Those two examples tell us that Google is not that worried about the cost of getting long-distance knowledge spillovers to happen within the firm. Evidently, they know that the knowledge will flow, and it doesn't really matter how far away the source is from the destination.

Could you discuss the home market effect and how it applies here?

Just to avoid confusion, the home market effect is a fundamental concept in trade theory that has nothing in particular to do with knowledge flows. But I find it a useful concept for thinking through the issues around technologies like AI. According to classical models of open economies, a country that has a lot of domestic demand for a product will, if trade is costly, produce more of that product: if domestic consumers want it, then local producers will make it. But as a country makes more and more of that product, it will see its production costs rise: producers will have to pay workers more, use more capital, more land, and eventually, as it tries to produce more and more of this product, in order to satisfy all these domestic consumers who want it, the country will become less and less efficient at making the thing. That also means that it will become less and less successful at selling the product abroad. However, if you think that we live in a world with strong local geographic increasing returns to scale, where industries learn from each other via things like knowledge spillovers, then everything is exactly the opposite of what I just described. In that case, a large domestic demand base allows the country to operate on a large scale at home. And it will become more efficient thanks to the fact that it is operating on that larger scale, which is, again, due to the fact that it has a large domestic demand base. Essentially, the home market effect is said to be true when a large domestic demand base helps you become a successful exporter country, as opposed to the opposite, when it hinders your export success.

MANY SERVICES THAT ARE IMPORTED OR EXPORTED WILL SOON BE CODIFIABLE

In relation to what you've just discussed, do you think that AI will help us become more efficient?

We don't know yet. It's like asking if AI will be successful at predicting the weather. The people who know the answer to that question are computer scientists and climatologists and statisticians. To take one example, think of supply chain management, which I think of as a purely computational problem. It's not economics, it's purely an optimization problem, one that some people say AI will make easier (and hence more efficient). I'm more interested in the question of what that improved production efficiency will do to the wider economy. Could it mean that firms are particularly better at doing complicated forms of economic production that involve many geographic steps and geographic units? Another question is what impact that will have on the way that trade is organized. There are two unknown factors at work here: the effect of AI on supply chain management, and the effect of supply chain management on the economy, be it the domestic economy or trade.

Can new technologies contribute to diversifying emerging countries' exports, such as knowledge-based services?

A simple model for thinking about the way AI is currently affecting the economy and production is to imagine that there's a task that needs to be done and there are two ways of doing the task. One way is to ask a human to do it and the other is to ask a computer to do it. The way I think about trade in services is like trade in tasks. We're often not talking about trade in finished goods but in some of the tasks of production, some of the steps involved in it. Of course, any task could also be a finished good, the final good that consumers buy. When I get a haircut, that's a service that's also a finished good. But if you think of services trade as trade in tasks, and consider a firm that makes goods, then it's as if there are now three ways to approach that firm's production. One is to use a human at home, another is to use a human who lives in another country, and the third is to use a computer that lives in a cloud somewhere. The way I think of a trade in services is that the technology for finding a human being domestically has always been available. That hasn't always been the case with the technology of finding a foreign human being who can do that task for you. We don't know how to do that for haircuts yet, but we do for radiographer services. Roughly 20 or 30 years ago it became cost effective to get a foreigner to perform those tasks, and my understanding is that this was due to a breakthrough in the ability to codify tasks and send communication about them over long distances. My sense is that the same features that make some tasks easy to codify and explain to a foreign human (and receive back from them) are the features that will make those tasks easy to instead codify and explain to a computer.

Does this mean that there will be trade diversion in these sectors? Offshoring?

Yes, absolutely. Maybe I'm wrong. Maybe there a lot of examples of tasks that are extremely hard for computers to do still but would be extremely easy to ask somebody in another country to do. But the examples I can think of are ones that you would have no hope of communicating to a computer if you couldn't first communicate them to another human being. It's as if all the things that we've figured out how to communicate to another human being over the internet are exactly the same things that would be easy to communicate to computers. If that were true, obviously that would suggest that AI would substitute for offshoring-type trade in services. But, again, that's just one example and I'm sure we can think of tasks for which that might go the other way.

Do you think that trade agreements should include some kind of regulation around AI, maybe to promote trade in services?

Trade agreements involve more than just tariffs, they involve regulations, too. A lot of standards and regulations act as protection, and as I've argued above, I don't think trade protection a good idea in most settings. When it comes to nontariff matters, I see a lot of trade agreements as just doing the good work of trying to reduce nontariff barriers to trade so that the world can trade more. That's a good thing and it should benefit people, on average, in very real ways. Finding ways to reduce nontariff measures so that countries can trade services more would be a wonderful thing on the whole.

Is there a trade-off between the right to privacy versus data as a driver of innovation?

The way I see it is that, in most cases, consumers sign an agreement with a digital platform before using it, and data use is part of the agreement they sign when doing business with a company. I don't think there's anything wrong with that. Maybe we need to educate consumers more so that they understand what they're doing when they sign agreements like that. In the case of the consumer finance industry, we think that educating people is helpful. Consumers sign agreements for very expensive credit cards with credit card firms that maybe they wouldn't sign if they were better educated, and that's been the motivation for all the initiatives around consumer finance protection. I see the situation with data use as being analogous. And so maybe there needs to be a bit more education. But my prediction is that's not going to change most consumers' behavior. Most consumers don't care that much, I think, about the privacy aspect of it. However, the leakage of private information is another thing entirely. I am no expert on the magnitude of this, but my own informal sense is that it does seem to be happening incredibly often. So perhaps the other thing that we need is harsher punishments for those who allow leakage to happen (just as we do when it comes to environmental damage, for example).

TECHNO-INTEGRATION

The Regional Public Good that has been created through the partnership between INTAL and Latinobarómetro has carried out over 20,000 exclusive surveys in 18 countries in the region to gauge Latin Americans' opinions on integration and new technologies.

24%
of Latin Americans think that artificial intelligence and robotics will create more jobs than they will destroy

23%
would be willing to travel in a car driven by a robot

22%
would be willing to undergo remote surgery

Caution or Fear of Change?

Willingness to travel in a vehicle driven by a robot, by country

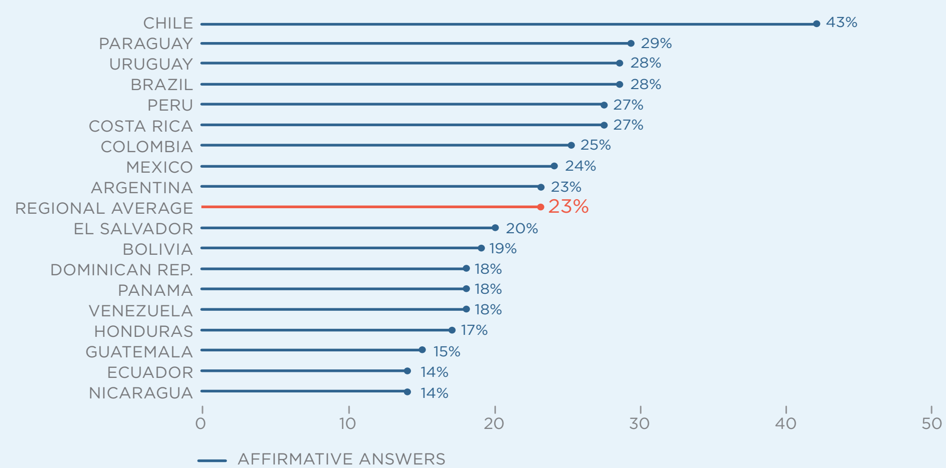

Source: INTAL-Latinobarómetro.

THE ABC OF ARTIFICIAL INTELLIGENCE

WHAT IS IT?

AUGMENTED INTELLIGENCE
Artificial intelligence is a technology that allows machines to carry out tasks as though humans are performing them. When used in combination with human intelligence it is referred to as augmented intelligence.

THE DATA DELUGE
Algorithms are ordered set of operations that solve a problem. They are the raw material for AI, which is now a factor of production.

BLACK BOXES
Self-learning and the automatic creation of new algorithms in overlapping layers generate results based on an underlying logic that is hard to track.

HOW IS IT DONE?

MACHINE LEARNING
Computers no longer need programmers: they learn rules by themselves based on experience (induction). In the case of deep learning, they are also capable of making their own decisions.

PATTERN RECOGNITION
AI tools can recognize images, speech, and emotions. There are multiple applications for this which touch on everything from real-time translation to the care economy. When these interact with one another, they form neural networks.

AUTONOMOUS SYSTEMS
These are devices that include the Internet of Things (IoT) and big data. They allow us to control objects and make decisions in small spaces. Driverless vehicles are based on this technology.

ROBOTICS
Robots are the most advanced form of AI. Today they operate in trade, industry, and services, carrying out complex tasks that only humans were capable of until very recently.

WHERE IS IT HAPPENING?

HEALTH
Robots now perform surgery and carry out medical diagnostic work based on image analysis using tools like IBM's Watson.

EDUCATION
AI allows us to personalize education and adapt teaching methods to each student's performance.

ECONOMICS
Algorithms play a part in financial transactions, digital marketing, and logistics using drones and autonomous vehicles. AI is changing every aspect of our economic life.

WHAT ARE THE CHALLENGES?

EMPLOYMENT
How can we handle the risk of job automation and prepare future generations for the labor market of tomorrow?

TRANSPARENCY
Algorithms make it essential for us to have access to data and decision-making formulas, which may be biased.

CYBERSECURITY
Respect for privacy and preventing hacking and data piracy are two areas that need to be strengthened if the digital economy is to work.

GOVERNANCE
We need an institutional design that is based on clear rules and ethical values that guarantee that the incorporation of AI into society will bring positive outcomes.

Transforming **Global** Services

Avinash Vashistha and Ankita Vashistha
Tholons Inc.

HOW WILL ARTIFICIAL INTELLIGENCE (AI) IMPACT GLOBAL SERVICES? THIS ARTICLE ANALYZES CASES SUCH AS MARKETING, SOCIAL AND LEGAL SERVICES, CYBERSECURITY, TRADE, AND OTHER AREAS WHERE NEW TECHNOLOGIES ARE BRINGING ABOUT DISRUPTIONS. IT EXAMINES THE MOST INNOVATIVE FIRMS AND THE START-UPS WITH THE GREATEST POTENTIAL IN EACH FIELD.

Over two decades of services globalization, opportunity has transformed many nations. India, Philippines, Eastern Europe, and Latin America have led the charge in shaping the history of this industry. Today, services globalization is worth over US$183 billion globally. The outsourcing industry has been shaken to the core and disrupted by digital forces. Robotics, artificial intelligence (AI), social media, mobility, big data, digital supply chains, digital trust, software-as-a-service, and cloud computing will continue to transform businesses and create opportunities for growth markets. There is an immense opportunity for spending on technology to grow from US$2.4 trillion to US$3.8 trillion, given the accelerated transition of legacy businesses to digital models. Digital innovation at the enterprise level can be achieved using the "five pillars" innovation strategy. Emerging countries that are not saddled with legacy businesses or major transformational efforts and costs would be justified in looking forward to a new dawn of certainty, opportunity, and innovation. The future of collaborative working is coworking; the future of innovation is co-innovation; and the future of investing is co-investing. These three collaborative elements are core to how enterprises will work, innovate, and invest in digital and the future globalization of services. A good foundation for countries is to implement the "open innovation" ecosystem to bring these elements together in a way that will also facilitate collaboration between growth leaders like India and Latin America.

CENTERS OF EXCELLENCE

The services globalization industry is now over two decades old. The industry traces its roots to the early 1990s. Countries like Ireland, the Philippines, India, Vietnam, Brazil, and the Czech Republic were the early destinations for taking information technology (IT) and business processes management (BPM) offshore. Very early on, these countries emerged as centers of excellence. Ireland was the pioneer in both IT and BPM. The Philippines emerged as a center of excellence for customer service; India for applications support and maintenance; Vietnam for telecommunications engineering; Czech Republic for BPM for European clients, for whom data privacy was important;

US$170 BILLION
THE POTENTIAL SIZE OF THE CYBER SECURITY MARKET

and Brazil for Latin American clients.

The industry has grown significantly over the last two decades and is now one of the largest sectors in terms of GDP, employment, and exports revenue for many countries. Data processing has evolved into BPM, knowledge-based services, and now e-commerce.

INDUSTRY SIZE AND OPPORTUNITIES

The global IT-BPM industry is now worth over US$183 billion, with India and Philippines being the top two offshore destinations. The United States and the United Kingdom continue to be the largest sourcing markets. The following data points summarize the global market size and opportunities: global sourcing reached US$183 billion in 2015 (8.9% growth) (NASSCOM, 2016a) and global spending on IT-BPM was US$1.25 trillion in 2015 (0.4% growth) (NASSCOM, 2016b).

Advances in technology are increasingly automating business processes while also reducing the cost-related benefits that were gained by outsourcing, which centered on taking advantage of low-cost infrastructure and talent. While digital disruption continues to be the prime mover in the IT-BPM industry, vertical markets like healthcare, finance and accounting, media, and government services have emerged as key drivers for growth.

Despite turmoil and concerns about the future of IT/BPM outsourcing, there are still significant opportunities for short-term growth. Increased connectivity between sourcing markets and providers—the outcome of both the proliferation of IT and the reduction of economic barriers—has created massive disruption and potential for growth and increased efficiency. In the next few years, the countries that are best able to take advantage of this will be in a position to take the infrastructure generated by the BPM boom and use it to their advantage in the future.

MAJOR OUTSOURCING LOCATIONS[1]

A large talent base for digital skills and the cost-effectiveness of labor continue to be the most important factors for outsourcing. However, there are other crucial factors that differentiate competitors. The reason that the Philippines is doing better than Latin America despite its relative size and that India continues to maintain such a dominant lead is English-language proficiency. The importance of English should have decreased with digital disruption; however, the results show the opposite.

India continues to be the top global outsourcing location, receiving 60% of the global outsourcing spend, followed by the Philippines, Latin America, China, and others. Its large talent base for digital skills and the cost-effectiveness of operating in the country remain the key factors in outsourcing. However, creative, digital, and innovation skills will define the future of jobs and industry revenues. This is an opportune time for countries like India and Latin America to collaborate and co-innovate to generate innovative industry solutions.

TRENDS SHAPING TODAY'S INDUSTRY

The emergence of digital technology has significantly impacted businesses by creating a mandate for them to adopt digital transformation. Major IT/BPO service providers are experiencing soaring growth in their digital portfolio, which in most cases is three to five times the average growth rate of the company. These companies should skill up and deploy business development and delivery capabilities across the following sectors:

Automation, robotics, AI, cognitive computing, and re-imagining the workforce. The workplace has changed. More and more tasks are being automated and a robotic workforce is emerging. Companies that reimagine their workforces and effectively blend humans and technology will establish a strong competitive position in the years ahead. Cognitive computing and AI are multipliers in the customer value delivery model. To fully realize the enormous potential of humans working together with technology, companies must emphasize retraining, helping people gain the skills needed to complement machine capabilities.

User Interface (UI) Design. Over the past few years, most of our day-to-day activities have shifted to the screens of our computers and mobile devices. Most businesses now have an online presence and workforces are being transformed to combine humans and robots, the UI design industry has the potential to grow to US$5.58 billion by 2019 (Transparency Market Research, 2016).

Digital Trust, Resilient Architecture, and Cyber Security. Discussion around IT security is now a much broader responsibility for all stakeholders. Today's technology needs to protect privacy, deliver benefits in exchange for the use of personal data, and demonstrate accountability. Strong digital trust will help brands attract and retain customers, deliver new products and services, and position brands well in the larger value chain for goods and services. Traditional defenses such as antivirus software and network firewalls have failed to stop the continuous stream of breaches. Cybersecurity and business risk management is now a broad agenda. The cybersecurity market is expected to grow to US$170 billion by 2020 (MarketsandMarkets, 2016a).

TABLE 1
THE FIVE PILLARS OF DIGITAL INNOVATION

CLIENT	START-UP PORTFOLIO	STRATEGY/ MENTORS	DIGITAL PARTNER	FUNDING
Critical mass of service providers and corporations with connections at the customer experience level	Critical mass of high-quality B2B start-ups will help build solutions and enable digital transformation	Provide strategy and consulting to start-ups and clients to align solutions and digital transformation	Systems integration/digital partner to implement, scale, and maintain solutions for large clients	Strategic partnership and investment in innovative companies

Source: Tholons Research 2017.

Big Data. Data is the new oil. It is the raw material for a digital economy. The high volumes, speeds, and varieties of data from companies' websites, social media, sensors, and so on have become a tool for businesses to get insight into customer's needs and behavior. Companies continuously look for data about their products, solutions, customers, and employees that can give them an edge over the competitors. Massive inefficiencies in processes can be eliminated using big data analysis. For FY 2015, analytics registered the highest growth in global spending in the BPM sector, and the big data market is expected to grow from US$28.65 billion in 2016 to US$66.79 billion by 2021 (MarketsandMarkets, 2016b).

The Data Supply Chain. Enterprises' data ecosystems have become complex and are in silos. This makes the data more difficult to access across the process, which in turn limits the value that organizations can get out of it. The challenge that we face today is efficiently linking and feeding that supply chain. There has been a lot of focus and innovation at the front and back ends of the data supply chain. Storage, which was once considered a prohibiting factor, has largely been solved thanks to solutions such as Amazon AWS. At the other end, there are hundreds if not thousands of companies creating algorithms and cognitive computing solutions that focus on business intelligence (BI) and AI and leverage big data to drive marketing.

Mobility/Apps. There is already a significant use of smart devices. Mobile subscriptions are growing at around 3% year-on-year globally and have reached 7.5 billion, while internet subscription rates have reached 5.1 billion (Ericsson, 2016). Devices across all industries are becoming mobile. These include the medical devices being used by patients and those used on manufacturing floors. This has given rise to the need for customized applications to be developed for each such function. Freelancers and SMEs would be well placed to develop business in mobile apps and testing.

Digital Marketing. With the exponential increase in mobile and internet penetration over the past few years, digital marketing has become the most promising marketing tool to effectively capture a wider customer base. Digital tools have enabled businesses to track and analyze customer behavior against any marketing campaign to increase its success rate and reach.

Content Management. Digital Content needs to be made available automatically, in real time, and at high velocity. Be it text, images, catalogues, audio, video, or interactive content, there is an opportunity to use platforms to be able to configure and deliver these products to clients.

The Internet of Things/Integrated Devices. The Internet of Things (IoT) is growing and has become significant, both in business and the personal world. Broadband Internet is becoming widely available and connectivity costs are decreasing significantly. There is a rapid increase in devices that are connected to the internet to enable the controlled operation and execution of task(s). These can be implanted across manufacturing floors in factory machines, in hazardous mines, oil rigs under ocean, race cars, and for surveillance systems. Knowledge of devices and the specific industrial process and the ability to develop embedded software will be a key niche expertise. These skills are in short supply, even in countries like the US, UK, Europe, Japan, and other places including outsourcing destinations like the Philippines, Eastern Europe, China, and India. Analysts have predicted that the installed base for devices will grow from around 10 billion connected devices today to as many as 30 billion devices by 2020, based on a 3 billion increase per year (Bauer, Patel, and Veira, 2014)

THE FIVE PILLARS OF DIGITAL INNOVATION

Enterprises are being challenged to integrate digital technology and innovation into their businesses. This is not just about global sourcing and outsourcing of IT and business processes. Innovation has to be a key part of the way enterprises work in the future. Enterprises need to collaborate with technology partners and start-ups to build innovative solutions. The future of innovation is co-innovation. Tholons' innovation strategy for enterprises is built on five pillars. These bring together clients, client partners, the network of global start-ups/innovation hubs, mentors, digital platform-builders, and investment to accelerate building enterprise solutions.

Connection with Clients. Connect with clients enables start-ups to not only have increased opportunities to market their products, they also get priceless engagement and feedback from the largest clients in their target industries.

Network of Start-ups. The network of global start-ups brings with it the best-of-breed start-ups to connect and collaborate. Furthermore, connections with other accelerators provide start-

5.1 BILLION
PEOPLE ARE INTERNET SUBSCRIBERS

ups with more places to seek business and funding.

Mentors/Partners. These are key to bridging the gap between enterprise clients and start-ups.

System IntegratorN/Digital Platform-Builders. These are critical to the process of integrating a solution from multiple start-ups and deploying it throughout a large enterprise and potentially across the globe.

Financing. Successful start-ups need funds to fuel their growth. Today, clients, technology partners, and venture funds co-invest together.

KNOWLEDGE-BASED SERVICES (KBS)

KBSs are programs that provide content-based (data, information, knowledge) organizational outputs (advice, answers, facilitation) to meet client needs. There are four types of KBS: generating content, developing products, providing assistance, and sharing solutions.

Transformation of KBS
Automation, robotics, analytics, and AI are shaping today's business processes. Intelligent automation is saving enterprises time and resources as well as increasing their efficiency. As intel-

ligent automation in the back office has evolved to provide virtual assistance on the front end, AI will become the digital face of brands (figure 3, see annex).

Finance and Accounting Services (F&A)

Robotics process automation (RPA) has transformed the F&A industry. This is a large segment of outsourcing and is already being automated through the use of software robots to replace human agents. Companies like BluePrism, UI Path, and Automate Anywhere are leading the RPA process.

Healthcare: Revenue Cycle Management

Health IT and electronic health record (EHR) systems have helped to streamline and improve the accuracy of healthcare revenue cycle management strategies. Many organizations use this technology to track claims throughout their lifecycles, collect payments, and address claim denials.

Legal Services

AI is now being integrated with the legal services market at an increasing pace. Robots now decide what paragraphs to include in legal contracts and help traditional lawyers, who are struggling to maintain the old order. Software has now moved on to smarter search and discovery, contracts, analysis, and more. Numerous AI platforms now create legal documents. Intelligent software reads, interprets, and extracts specific information from documents and converts it into a desired structured output, in a fraction of the time it would take humans to do so.

A few of the large global audit firms, including KPMG, are feeding audit data to IBM's Watson, which provides comprehensive, unparalleled audit support for their teams. Predictive analytics allows for the automation of evidence gathering and the production of complex data reports, saving time and improving client services. Very soon, these capabilities will be far superior to what can be done by experienced auditors.

Cognitive Computing & Customer Relationship Management

Digital marketing bots and AIs are just scripts that are integrated into an app interface. They rapidly scan the user's text input and deliver a response that has been predesigned and preloaded by a developer. So far, bots and AIs have typically performed a customer service role, answering queries and directing users to helpful information and resources. Now they are also being used to participate in or even initiate sales conversations. The biggest existing platform for sales-focused implementation of bots and AIs is the Facebook Messenger app, which announced the launch of a "sales bot development" ecosystem for ad clients in April 2016.

Another good example would be Amelia, by IPsoft, a virtual agent that functions as a call center operator. Amelia can be deployed straight from the cloud in a fraction of the time it would take to train a human. This AI learns as it works and provides high-quality responses in various languages. The new development on Amelia has improved her memory, contextual comprehension, and emotional responsiveness. This project is getting close to passing the Turing test.

Digital Marketing

AI is not only being used in digital marketing, it has already started making the lives of users and marketers easier. From texting to visualizing business insights, the merging of big data, machine learning, and AI is creating smoother and smarter experiences every day. Modern marketing and digital marketing are characterized by the rapid use of consumer data. AI technologies are simply far superior to humans when it comes to processing and understanding vast data sets, and producing evidence-supported data insights. Automation will be the output of AI.

Data Analytics

Data analytics and insights are helping business processes by providing deep insights into customer behavior, market trends, risk management, decision-making, increasing productivity, and work-flow management. In the world of healthcare, IBM Watson allows professionals to spend less time searching literature and the EMR and more time caring for patients. Watson can provide clinicians with evidence-based treatment options based on expert training by Memorial Sloan Kettering (MSK) physicians for cancer treatment.

Market/Business Research

Big data, social analytics, and text mining are new ways of carrying out market research. Research can be easily scaled to the international level, and working with text and images has become easier. Data collection devices and tools have increased the pace of data creation. With real-time data capturing it has become easier to obtain valuable insights into data. Big data provides gives tremendous opportunities for researchers to gain real insights from huge data sets.

SERVICES IN E-COMMERCE

E-commerce companies are spending huge amounts of money on improving the online shopping experience and increasing sales. Technologies like digital payment gateways, cataloging, virtual and augmented reality, extensive use of analytics for consumer and sales insights have changed, or are set to change, the face of e-commerce. Services in e-commerce are constantly growing.

Transformation of E-commerce Services

Cataloging. Enhancements such as innovative cataloging, interactive product displays, and dynamic personalization are enabling retailers to deliver an online experience that is more aligned with consumer preferences.

Digital marketing. Developments in digital marketing such as personalized retargeting, triggered communication, SMS/push messages, and master data management are modern ways of consumer targeting.

Virtual/Augmented Reality: The IKEA mobile app includes an augmented reality feature to give consumers a virtual preview of furniture in a room, allowing for a digital test-run of the brand's products.

Logistics. Subscription boxes can reduce violent peaks and troughs, providing more predictability and fewer instances of products being out-of-stocks and backlogs caused by demand-spikes.

Digital Payment. Digital payment systems have simplified accounting, improved efficiency and security, and reduced administration costs. Many banking services are being redefined. This includes technologies related to e-commerce, mobile payments, currency conversion, etc.

Data Analytics. The biggest impact on shipping and fulfillment in 2017 was the use of analytics and AI. Chatbots are a great example of how AI technology is becoming more widespread and more helpful. A smart chatbot might monitor your online buying habits and make a purchase on your behalf when it sees that a product you've bought before and reviewed positively goes on sale.

The Data Supply Chain. The use of big data and analytics have made it possible to make real-time deliveries, improve vendor management, automate product sourcing, and personalize segmented supply chains.

THE DIGITAL TRANSFORMATION OF INDUSTRIES

Businesses must embrace digital technologies and platforms like social media, mobile technology, cognitive computing, cloud computing, and AI. Digital transformation is shaping the customer experience, operational processes, and the business model. Digital capabilities have become much more important within any organization. To make their digital journey successful, industries have started investing in digital initiatives and skills.

Demand for Digital Innovation

Digital technologies are becoming an essential enabler to more sustainable lifestyles, business models, and societal institutions. They are opening up new options in design, manufacturing, IoT, health, banking and the financial sector, retail, biotechnology, and so on.

Digital innovation in healthcare is improving and advancing healthcare processes, decision-making, and patient experiences. Players in the healthcare industry are constantly moving toward IT adoption. Figure 4 is an example of how innovative companies are unbundling legacy healthcare to make it more seamless, effective, and affordable.

Chronic disease management is currently complex, expensive, and of questionable quality. New technology and processes using IoT devices to monitor, diagnose, and guide treatment are changing the way that chronic diseases like diabetes, cardiac ailments, and oncology are being managed. Figure 5 gives examples of these solutions for diabetes and cardiology.

A new wave of technologies—such as digital wallets, robotics, and artificial intelligence—is revolutionizing the way we manage, control, and distribute money. Almost every area of the finance and banking industry is being disrupted and forcing traditional banking to innovate and invest in digital innovation. Figure 6 gives an overview of how financial services are being disrupted by new technology start-ups.

Retail

Digital innovation is improving online and in-store customer experiences. Technologies like virtual reality, digital shelves, payment through facial recognition, click and collect services, and mobile coupons are adding significant value to shopping. Figure 7 captures some of the innovative start-ups challenging the traditional retail business model.

Manufacturing and IoT. IoT is changing industry dynamics and bringing operational innovation and excellence to smart manufacturing and the factories of the future. Figure 8 highlights some of the areas where IoT is being incorporated and various start-ups working in these areas.

Energy and Utilities. Digital innovations are making a significant impact on sustainability and the use of renewable energy. Technology provides utility companies and consumers with ways of controlling their consumption, thus reducing usage and the cost of energy. Figure 9 shows technology start-ups that are transforming the energy and utility sectors.

Biotechnology, Genomics, and Pharmaceuticals Biotech, genomics, and pharmaceutical companies are bringing digital revolution to healthcare and clinical research. Digital technology has increased patient engagement and most care has been protocolized, which has facilitated clinical decision-making.

Media and Entertainment. The digitalization of the media and entertainment industry has changed consumer behavior and expectations. We can now access content anytime and anywhere. Technologies like virtual reality, high definition graphics, and animation have brought consumers new experiences.

Virtual Reality (VR) and Augmented Reality (AR). VR and AR are two powerful technologies that can change the experience of creating, selling, and buying products. They are also being integrated heavily in the gaming and entertainment industries.

Cybersecurity. Digital innovation needs robust digital security, without which system data will be at constant risk from hackers and other digital threats. To successfully compete and maintain a solid position in the market, a business needs to have a good cybersecurity system. When different apps communicate and share data, digital trust becomes important. Systems need to build active defense mechanisms, resilient architecture, and robust authentication procedures. Figure 10 gives an overview of innovative start-ups in various areas of cybersecurity.

COMPETITIVENESS IN SERVICES GLOBALIZATION

Tholons Services Globalization Index, published annually, is the industry's premier rating and ranking of the Top 100 Super Cities and Top 50 Digital Nations. The key elements defining the ranking include infrastructure, cost, talent, business maturity, risk, quality of life, and digitalization/innovation. Countries' and cities' competitiveness is shifting as a result of digital innovation and transformation.

Digital technology is now a critical element in disrupting and transforming industries globally. New technologies, business process management companies, and multinational corporations need to align with the stark reality of digital disruption. For the front-runners in services globalization—India, the Philippines, and Eastern Europe—most of their services will be commoditized. The digital transformation of clients and services such as mobility, analytics, and cybersecurity cannot be delivered using traditional business models. Moreover, the ongoing automation of work—in the form of application management, infrastructure support, and testing—is reducing or in some cases eliminating the need for manpower. These forces will severely affect IT service companies.

Tholons Services Globalization 2017 includes innovation, start-up ecosystems, and digital transformation as key components in its Digital Nations (Top

50) and Super Cities (Top 100) indexes. Digital technology is changing the industry and shaping the leaders, disruptors, and innovators that will one day define the future of economies and growth markets.

The outsourcing industry has been shaken to the core, and major industry leaders are slipping from top positions. Pune (India) and Cebu City (The Philippines) have moved out of the top 10 while São Paulo (Brazil) and Buenos Aires (Argentina) have moved into the top 10, showing the significant inroads being made by Latin American cities. Similarly, Canada, Chile, and Ireland have moved into the top 10 on the Digital Nations ranking.

Disruptors and innovators on both these lists will continue to challenge the established leaders. Firms that engage exclusively in e-commerce in these digital nations and super cities will define the future of services globalization.

OPEN INNOVATION ECOSYSTEM

Open innovation revolves around building an ecosystem wherein start-ups, academia (universities), private institutions (IT/BPM companies, technology suppliers) and financial institutions (venture capitalists, banks, angel investors, and so on) come together to foster collaboration and co-innovation. Start-ups are expected to conceptualize and develop products or services, universities are expected to generate quality graduates to feed into start-ups, and financial institutions are expected to provide the funding for the start-ups to develop their products and services.

The open innovation ecosystem needs support from government through policies and incentives that are aligned to promote innovation, the entrepreneurial spirit, and collaborative work between all stakeholders.

Open innovation is key in today's world where innovative products and services are disrupting the major industries. The start-up ecosystem is particularly vibrant in the US, the UK, Israel, India, and Canada. There are 100,000+ start-ups in the US, 7,500+ in the UK, 5,000+ in Israel, 7,000+ in India, and 5,000+ in the Middle East. Developing a robust Open Innovation Platform with industry stakeholders and government agencies should be considered when developing a country's IT/BPM ecosystem.

Countries would be well served by implementing the open innovation ecosystem and creating a start-up culture that is connected to the global network of clients, start-ups, mentors, platform-builders, and funding. The following is a guide for starting this process:

Step 1. Creating awareness among start-ups. Awareness needs to be created among entrepreneurs by promoting start-up fund events on social media and at various start-up hubs. These provide selected start-ups with significant visibility, connecting them to a network of mentors and giving them access to investors for funding.

Step 2. Screening and due diligence. After the successful registration of applications from various start-ups, a due diligence process selects 30–50 start-ups that will go through a selection and mentoring process. Selection is based on basic fundability elements: idea potential; the size and scalability of the opportunity; the stage of the company's development; team profile; and investor attractiveness.

Step 3. Preparing for the pitch session. Selected start-ups will receive instruction on expectations around their pitches and other guidelines. Start-ups will be expected to submit pitch videos and pitch presentations. They will receive mentoring and advice on the best ways to pitch and approach investors.

Step 4. Pitch session. This is a day-long event with live presentations to 5–10 jury members and investors. Each pitch session lasts anywhere from 10 to 30 minutes. In the end, the jury and investors select start-ups to take on board for further mentoring and investment. Three to five start-ups are selected and sent to spend three months at a Silicon Valley tech hub where they can cowork and co-innovate with multiple other start-ups and clients.

COWORKING, CO-INNOVATION, AND CO-INVESTMENT

The future of working is coworking; the future of innovation is co-innovation; and the future of investment is co-investment. These three collaborative elements are at the core of how enterprises will operate, innovate, and invest in the digital transformation of industries and consumer experience. The following three sections define these elements, discuss the trends and benefits, and outline how they can be developed and implemented.

FIGURE 1
THOLONS SERVICES GLOBALIZATION INDEX 2017
TOP 50 DIGITAL NATIONS

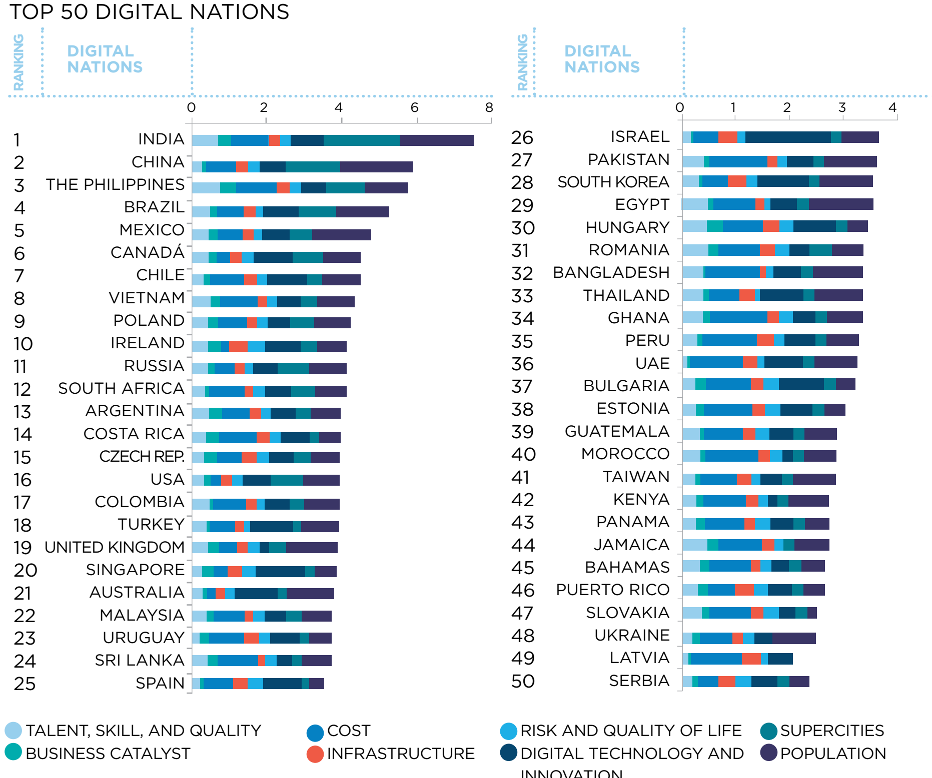

Source: Tholons Research 2017.

Coworking

Coworking is about more than just affordable shared workspaces. It is about working in a community and sharing ideas and helping each other. Freelancers, entrepreneurs, and start-ups can start to feel part of the community, build connections, and pay to work in a communal space. The economic benefits from coworking are significant. According to Tholons research data, on an average, coworking brings savings of as much as 55% compared to operating an individually owned workspace.

Coworking is a need and an opportunity. It is needed by the 1.6 million people who subscribed to a coworking space this year and it is an opportunity for dynamic individuals to connect with one other.

Coworking is more than just a shared working environment. It is about building a community. Nearly 34% of the US workforce operate as freelancers. Growing numbers of entrepreneurs and start-ups have led to the demand for coworking spaces doubling. As many as 65 million small businesses will join coworking spaces by 2020. Of these, 2.5 million will do so by 2018. Demand for coworking spaces currently outstrips supply by a factor of 3:1.

Community is the lifeblood of coworking and co-innovation. Workshops, mentor talks, pitch sessions, demo days, and networking events all make the community meaningful.

FIGURE 2
THOLONS SERVICES GLOBALIZATION INDEX 2017
TOP 50 SUPER CITIES

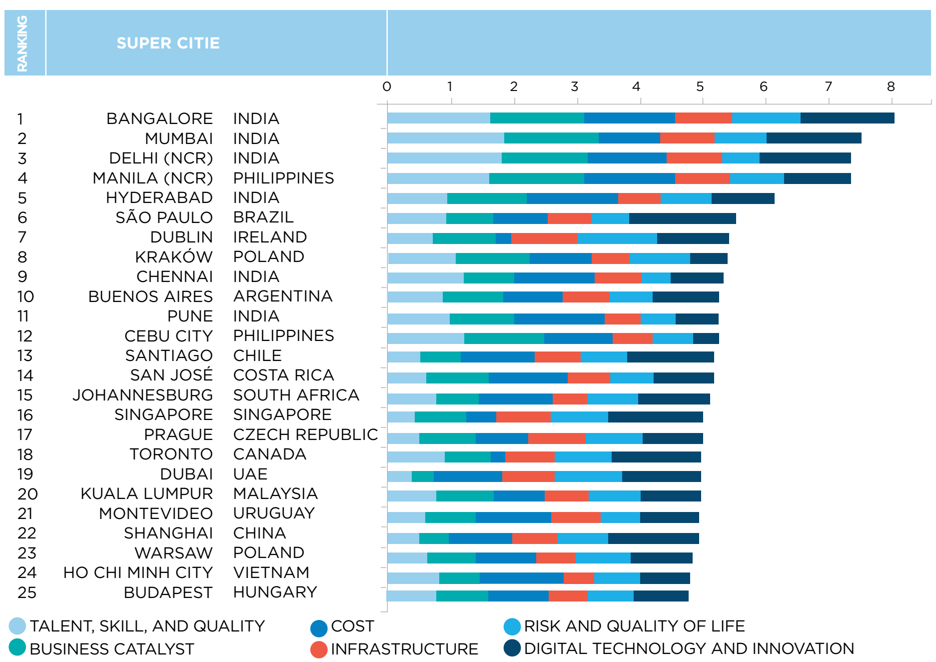

Source: Tholons Research 2017.

Co-innovation

Businesses are becoming more and more connected. Individual efforts alone are not sufficient for innovation. In the world of innovation, collaboration is becoming imperative. Co-innovation involves two or more partners that deliberately manage mutual knowledge flows across their organizational boundaries through joint invention and commercialization activities. Equity sharing between clients, mentors, investors, and strategic integrators provides a robust support system for start-ups.

Co-investing

The revolutionary involvement of partners in the co-innovation model relies on an equally pioneering yet simple cofunding strategy. Co-innovation is fueled by clients, technology partners, and venture capitalists investing alongside each other.

Digital Value Chains

The services globalization industry is transforming many emerging markets, including India, the Philippines, Latin America, and Eastern Europe. IT, BPM, KPO, and e-commerce are now a US$183 billion industry. Services globalization has been shaken up by digital forces. Robotics, AI, social media, mobility, big data, digital supply chains, digital trust, and software-as-a-service will continue to unravel established businesses and create opportunities for many countries. The outsourcing industry has been shaken to the core, and major industry leaders are slipping from top positions.

Spending on technology currently stands at US$2.4 trillion and is slated to increase to US$3.8 trillion through the accelerated transition of legacy businesses to digital ones. This opportunity should not be missed. It is good news for innovative start-ups, super cities, and digital nations, who are in a position to stake a claim in the new digital landscape. Latin American countries and cities have shown some of the most remarkable growth in this area. Given India's leadership and the growth of Latin American countries, there is a good business case for these two powers to collaborate and co-innovate. A commitment to focusing energy and resources at the enterprise level will lead to the emergence of new leaders. These new players are in an enviable position as they are not saddled with legacy structures and the effort and cost that transforming these entail. Emerging countries should be keen to embrace the opportunities that this newfound world of digital technology and innovation holds for them.

NOTES
[1] NASSCOM (2016a).
[2] NASSCOM (2016b).
[3] NASSCOM (2016b).

REFERENCES
Bauer, H., Patel, M., and Veira, J. 2014. The Internet of Things: Sizing up the Opportunity. New York: McKinsey & Company.
Ericsson. 2016. Ericsson Mobility Report: On the Pulse of the Networked Society. Stockholm: Ericsson.
MarketsandMarkets. 2016a. Cyber Security Market by Solution 2016. Pune: MarketsandMarkets Inc.
—. 2016b. Big Data Market by Component Report 2016. Seattle: MarketsandMarkets Inc.
Nasscom. 2016a. NASSCOM Strategic Review 2016. New Delhi: Nasscom.
—. 2016b. Strategic Review 2016: The IT-BPM Sector in India. Tholons Research 2016.
Transparency Market Research. 2016. Human Machine Interface Market 2016. Albany, NY: Transparency Market Research.

APPENDIX: THE DIGITAL UNIVERSE

FIGURE 3
ARTIFICIAL INTELLIGENCE

FIGURE 4
DIGITAL HEALTHCARE

FIGURE 5
DIGITAL CARE

FIGURE 6
FINTECH

FIGURE 7
RETAIL TRADE

FIGURE 8
MANUFACTURING AND IOT

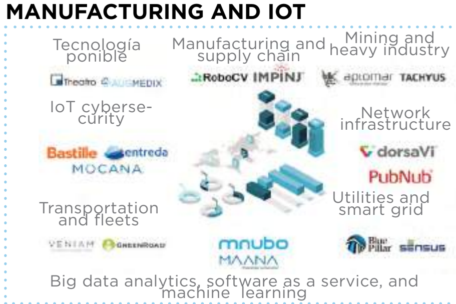

FIGURE 9
ENERGY AND UTILITIES

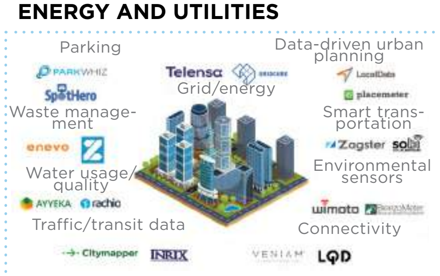

FIGURE 10
CYBERSECURITY

Source: Tholons Research 2017.

Capacity Transfers

Synergies between Academia and the Business World

Ricardo Oscar Rodríguez
UBA/CONICET

Argentina's scientific community includes just over 100 professional researchers and about the same number of trainees in the field of artificial intelligence (AI). The local AI ecosystem has given rise not only to internationally recognized theoretical developments but also to multiple innovative applications.

At the National University of the South in Bahía Blanca, researchers are working on e-governance through a project called DECIDE 2.0, in partnership with academics at the Monterrey Institute of Technology and Higher Education in Mexico (ITESM) and the United Nations University International Institute for Software Technology (UNU-IIST) in Macau, China. The aim of the project is to develop an intelligent processing framework for public opinion as expressed on social media (such as Facebook, Twitter, Instagram, Snapchat, Waze, etc.), using a collaborative system that operates on top of these networks. Through this, the team is seeking to implement models based on trust and reputation propagation so that decision-makers can evaluate people's opinions of others' reputations appropriately. Furthermore, as different users can give their opinions on a topic, the project also aims to develop algorithms to include information from different sources. Likewise, it seeks to identify new topical issues by obtaining conceptual information associated with these, with the eventual aim of developing models for specific fields (transportation, health, education, security, etc.). The project brings together many of the theoretical results that researchers in the group have been working on over the last decade, such as the analysis of feelings, data mining, lines of argument, visualization, and so on.

At the University of Buenos Aires's School of Exact Sciences, there are several projects and groups that transfer information, although I will limit myself to describing just two of them here. The first is the drone development project being carried out by the Laboratory of Robotics and Embedded Systems in collaboration with the Institute of Automation at the University of San Juan. What is unusual about these drones is the use of an innovative autonomous navigation system that improves their flexibility and response speeds when unforeseen circumstances arise. These drones have been used for virtual forest management and to implement precision agriculture.

The other development came from the Laboratory of Applied Artificial Intelligence. The team there have developed a tool that records speech, analyzes it, and detects whether the speaker may suffer from schizophrenia based on their speech patterns (quantity of verbs used, disorganization, or discursive coherence). This application was developed in collaboration with North American psychiatrists and is based on the text analysis techniques this research group had developed.

Another research group based at the School of Mathematics, Astronomy, Physics, and Computing at the National University of Córdoba specializes in natural

100+ RESEARCHERS ARE WORKING ON AI WITHIN ARGENTINA'S SCIENTIFIC ECOSYSTEM

language processing. They have been working on the MIREL project, which is partially financed by the European Union's Horizon 2020 program, which centers on developing analysis and reasoning tools for use with legal texts. One interesting feature of this project is that it involves European companies who are interested in using these tools to provide legal advisory services. It seeks to generate intelligent techniques for maintaining, accessing, and checking large repositories of legal data and to carry out reasoning tasks in relation to these. The aim is for these intelligent systems to totally or partially perform jobs such as searching for precedents and case law, providing counsel, and making decisions and so on. The systems in question use deontic reasoning models, ontologies, semantic text processing, and data mining.

The Research Institute for Signals, Systems, and Computational Intelligence at the National University of the Litoral, sinc(i), has patented methods, processes, and devices that use AI techniques for use in cattle raising or farming. For example, these can calculate when cows are in heat and can detect, classify, and quantify ruminants' feeding activity in real time. What sets this particular group apart is the large number of technologies it has developed based on signal processing, using devices such as the Holter HT103 or others that measure respiratory rate control and so on. They have also developed various biotech applications using data mining.

In Mendoza, the Laboratory for the Development of Tools for Machine Learning and Reasoning (DHARMa) is developing the Vineyard Sunlight Exposure (ViSE) project, which automates the creation of 3D models of grapevines and correctly identifies their components. The system aims to provide a low-cost tool for monitoring grape crops under real conditions and for implementing corrective action to improve growth. The techniques developed have proven stable enough to be used to detect, recognize, and identify the shape, size, and state of all parts of the plants.

A team at the Franco-Argentinian International Center for Data Science and Systems (CIFASIS) in Rosario has built a weeding robot that identifies different weed species in real time, which allows it to then treat each one appropriately. This mobile platform moves through fields autonomously using machine learning and artificial vision techniques.

The Systems Institute of Tandil at the National University of Central Buenos Aires has developed a service platform to produce software for smart cities and is committed to providing top-notch services and developing software applications for these. Its platform-based strategy seeks to integrate systems and data sources that were not necessarily designed for this end and to provide smart services for different applications. The first prototype will be used in Tandil.

As can be seen, the world of academia is not just about abstract issues: it also focuses on developing the technological capacities needed for transferring essential knowledge to the productive sector so that it can access cutting-edge technology, improve the quality of its products, and increase value chains. Outside of these academic spaces, there are also many entrepreneurs who are using AI techniques to design competitive computing applications.

Georeferencing and Big Data

A Trade Model Using Spatial Econometrics

Marcos Herrera
CONICET-IELDE, National University of Salta

THE GROWING AVAILABILITY OF GEOREFERENCED DATA HAS LED TO A NEED FOR SPECIFIC TOOLS TO CAPTURE THIS INFORMATION. ECONOMIC PREDICTIONS CAN BE ENHANCED THROUGH MACHINE LEARNING OR SATELLITE IMAGE RECOGNITION. SPATIAL ECONOMETRICS ENABLE INTERACTION EFFECTS THAT ORIGINATE IN GEOGRAPHIC SPACES TO BE DETECTED, AS CAN BE SHOWN USING A GRAVITATIONAL TRADE MODEL.

The growing availability of georeferenced data from individuals and different statistical agencies from around the world has led to the growing use of analytical techniques and spatial econometrics. Spatial econometrics is a branch of econometrics that analyzes *spatial effects* on regression models. These *spatial effects* (dependence and heterogeneity) come from the proximity or closeness between units of observation. This proximity may be geographical (closeness between geographical positions), social (closeness in terms of family or social ties), economic (closeness in terms of trade interactions or sector relationships), or combinations of these. Regardless of which of these proximity relationships applies, spatial econometrics considers the impact of these interactions to be especially relevant and generates models that allow you to estimate the effect of these.

Although these might seem to be recent developments, they actually date back to 1979 (Anselin, 2010), which was the starting point for different methodological developments that eventually led to the book *Spatial Econometrics: Methods and Models* (Anselin, 1988). In the early years, spatial econometrics only spread through the field in a limited fashion due to the scarcity of this kind of data and the shortage of statistical estimation programs. Both of these barriers have now been overcome, particularly due to the increased availability of georeferenced data from individual users, companies that allow access to their data via an application programming interface (API), and through programming techniques that extract data from websites (web scraping). Spatial estimation methods are also included in the most popular statistical programs, such as Stata, R, Matlab, and Python, among others.

Arribas-Bel (2014) stresses that the quantity and diversity of these new sources of information are accidental, open, and available everywhere on earth. The data is accidental in that it was not obtained through a survey or census that was specifically created for research or analysis of social or economic policies. Furthermore, it is freely available to researchers. The sources for this data are so diverse and on such different scales that they allow researchers to reduce localization error in observations, avoid the discretization of continuous problems in space, and draw on information that lies beyond data collected in traditional ways (surveys and censuses). But conventional statistical and econometric techniques are not yet ready to handle today's volumes of data, which is why *big data* techniques have emerged.

Likewise, they are unable to appropriately capture complex structures like interconnectivity and the geotagging of observations, which spatial econometrics is able to do.

What makes a spatial econometrics model different is the presence of a contacts matrix or connectivity network. This connectivity network can be seen as overlapping with the literature on social networks. The two fields differ, however, in that the focus of attention in social network studies is on analyzing the structure of connectivity, which is a way of characterizing each network using different measures (centrality, clustering coefficient, average score, network density, etc.), while spatial econometrics focuses on the impact of the connectivity network on the econometric model. These two fields of research clearly have points in common and may complement each other.

There is also an overlap with *big data*, in the form of geotagged data (*geo big data*), which provides alternative strategies for capturing spatial effects. For example, the *Livehoods* project seeks to define urban areas not only in terms of the places they are geographically close to but also through the people whose daily routines bring them to the area. Using over 18 million *Foursquare* records, the project groups places into areas according to patterns from the groups recorded in them. These check-ins (ratings and opinions of the places in question) provide information about different parts of a given city, which allows researchers to study the social dynamics, structure, and characteristics of cities on a large scale. These types of projects construct digital neighborhoods by combining user preferences and georeferenced information.

As can be seen, there are many different approaches to spatial data analysis and the development of these is in full swing. In the rest of this article, I will attempt to highlight the main spatial econometrics tools used for capturing complex information. I should begin by noting that spatial data has very specific features, regardless of whether the data comes from a traditional database or one of the huge databases used in *big data analysis*.

THE NATURE OF SPATIAL DATA

The first point I wish to stress is that there are different types of spatial data, which has led to particular statistical approaches to each. Cressie (2015) created a data taxonomy based on assumptions about the stochastic spatial process (*spatial random field models*). Without going into too much detail, spatial data can be classified as *geostatistical data*), regional or *lattice data*, and *point pattern data*. Likewise, this data can be represented and visualized in vector or raster format, which are two types of spatial information layer commonly used in geographic information systems (GIS).

In raster format, space is divided into regular cells and each cell contains a number that identifies the object. This style is recommended for representing geostatistical data such as surface temperatures, pollution, precipitation, etc. Satellite images from remote sensors (satellites or drones) are stored using this format.

In vector format, different objects are represented by points, lines, and polygons. This style is typically used for representing regional data, such as a map of a country divided into provinces or municipalities, or railway lines and locations of houses that are for sale. This is the type of georeferenced data that is used most in the social sciences in general and economics in particular.

But the two formats can be combined, and it is not unusual today to find mixed applications. For example, Patino and Duque (2013) review regional science applications that use satellite imagery of informal urban settlements. The most common applications use images from Landsat, SPOT, and ASTER, among others, to detect depressed urban areas and assess quality of life, urban growth, and social vulnerability. In a landmark study, Jean et al. (2016) analyze satellite imagery, vector data, and official survey data using machine learning techniques to estimate consumption and wealth in five African countries. By combining data and these techniques, the authors were able to identify poor areas that could be targeted by specific policies.

One common drawback to all approaches to spatial data is the *modifiable area unit problem* (MAUP), a geographical version of the ecological fallacy, where conclusions based on a particular aggregation of areas or regions may change if the same data is aggregated into a different set of areas or regions. This affects the inferential results and is not always highlighted sufficiently. Furthermore, the use of diverse data sources also brings risks—for example, data is often not available on the desired scale. When spatial data needs to be transformed to the required scale, the *change of support problem* (COSP) arises. The term “*support”* refers to the size and volume of each database, but also includes the form, size, and spatial orientation of the objects or fields that are represented. Using a different support generates new variables that are related to the original but are spatially and statistically different.

The particular qualities of spatial data are often not taken into consideration when researchers work with it. One such example would be the gravity model of trade, in which countries (polygons) are transformed into points (generally the capital city) to measure distances and these measurements are used as a variable. This simple transformation may entail all the problems discussed above. In this case, the exactness of the measurement may not significantly affect the results. However, progress in product geotagging may soon prompt the need for greater precision in the origin and destination of each good and may force us to rethink issues as simple as how to measure this variable.

ECONOMETRIC MODELS

Whether the space in question is geographical, digital, or socio-economic, the tools developed by spatial econometrics allow as to obtain a measure for the interaction and significance of spatial or network effects. Their main virtue lies in acknowledging the dependent nature of data from the outset, which is not true of standard statistical methods.

Standard econometrics makes observations independent from one another and thus makes the handling of statistics manageable. One implication of this assumption is that y_i values observed for individual i are statistically independent of y_i values for individual j, such that $E(y_i y_j)=E(y_i)E(y_j)=0$, where $E(\text{-})$ is the expected operator, assuming that values are centered, for simplicity’s sake. However, when individuals interact in such a way that their decisions are dependent on their peers or neighbors, this assumption must be relaxed to allow for some form of dependence.

The traditional form of spatial dependence in a cross-section is through spatial autocorrelation, which is understood as the similarity between values in nearby locations. Autocorrelation can be positive: a high (low) value for a random variable in a spatial position has a neighboring location with high (low) values. Negative spatial autocorrelation is also possible: a high (low) value for one geographical position comes with low (high) values for a neighboring one. The presence of spatial autocorrelation implies that a sample of autocorrelated data contains less information than an uncorrelated one. This loss of information needs to be considered explicitly in estimations and is the main problem in the use of applied econometrics using spatial data.

SPATIAL WEIGHTS MATRIX

Spatial autocorrelation can be understood as a rescaled version of covariance: $Cov(y_i,y_j)=E(y_iy_j)-E(y_i)E(y_j)\neq 0, \forall\, i\neq j$, where y_i, y_j are observations of a random variable at location i and j in space. In other words, each pair (i, j) contains specific geographic information measured by latitude and longitude. However, for a sample of n observations in a cross-section, there are $(n^2-n)/2$ covariances because they are estimated under symmetry and there is not enough data for all the pairs to be estimated. The way to solve this problem is by placing restrictions on the way in which observations interact with one another.

The main approach used in spatial econometrics for imposing restrictions on interaction is the spatial weights matrix, which is commonly referred to as "W" (or "G" in social networks) and describes the connectivity between n units that are located in a two-dimensional space. The construction of W is based on at least two key assumptions about spatial structure: (1) a connectivity criterion that defines which units can be considered to be neighboring one other and (2) a spatial weight assumption that operationalizes how neighbors affect each other.

To understand how the W matrix allows us to simplify the problem of spatial dependence, I will formalize the discussion by considering a *spatial autoregressive* (SAR) in which the variable y is spatially distributed in three regions, as follows:

$$y_i=\alpha_{ij}\, y_j+\alpha_{ik}\, y_k+u_i,$$
$$y_j=\alpha_{ji}\, y_i+\alpha_{jk}\, y_k+u_j,$$
$$y_k=\alpha_{ki}\, y_i+\alpha_{kj}\, y_j+u_k,$$

where $u_i; u_j; u_k \sim i.i.d.(0;\sigma^2)$ In other words, in the first equation, value y in region i depends on the value of y in region j and k, in addition to a random variable that is distributed identically and independently between locations. The same can be said of regions j and k In matrix terms, the system can be restated as:

$$\begin{bmatrix} y_i \\ y_j \\ y_k \end{bmatrix} = \begin{bmatrix} 0 & \alpha_{ij} & \alpha_{ik} \\ \alpha_{ji} & 0 & \alpha_{jk} \\ \alpha_{ki} & \alpha_{kj} & 0 \end{bmatrix} \begin{bmatrix} y_i \\ y_j \\ y_k \end{bmatrix} + \begin{bmatrix} u_i \\ u_j \\ u_k \end{bmatrix}$$

$y=Ay+u$
where

$$A=\begin{bmatrix} 0 & \alpha_{ij} & \alpha_{ik} \\ \alpha_{ji} & 0 & \alpha_{jk} \\ \alpha_{ki} & \alpha_{kj} & 0 \end{bmatrix}$$

The problem with this system is that there are more parameters and observations in a cross-section (three observations on seven parameters, including the dispersion parameter). Symmetry could be imposed by reducing the number of parameters that cannot be estimated all the same. Consequently, we need to find an alternative solution to be able to estimate the underlying dependence in the data.

The solution lies in imposing a set of restrictions on the dependent relationships. In this way, the structure of matrix A is reparameterized as follows:

$A=\rho W$

where ρ is a parameter to be estimated that captures the average effect of interaction between all the regions and W is the spatial weights matrix. The components of W are:

$$W=\begin{bmatrix} 0 & W_{ij} & W_{ik} \\ W_{ji} & 0 & W_{jk} \\ W_{ki} & W_{kj} & 0 \end{bmatrix}$$

such that the parameters α (the model's original parameters) have been replaced with W (coefficients that are exogenous to the model) and the model is expressed as $y=\rho Wy+u$, where Wy is interpreted as the spatial lag of y by analogy to time series.

The W matrix will generally be of the order $n \times n$, where n is the sample size. Each element of W is described as a spatial weight, W_{ij}. Spatial weights capture whether elements are neighboring and are different to zero when the i and j are considered to neighbor one another. By convention, no region can neighbor itself, which results in the main diagonal of W containing all those elements that are equal to zero, $W_{ii}=0$.

The most traditional way of building the matrix is using geographical criteria, in line with the first law of geography (Tobler, 1970): "everything is related to everything else, but near things are more related than distant things." Within geographic criteria, "neighboring" can be defined by adjacency, by using some function of distance, by *k* nearest neighbors, or by some combination of the above.

STANDARD SPECIFICATIONS

One way of estimating an econometric model with spatial data is via a nonspatial model, usually estimated through ordinary least squares (OLS). Using the residuals of the nonspatial model, the possible presence of dependence between observations can be contrasted using Moran's I test, the null hypothesis of which is no spatial autocorrelation. If the null hypothesis is rejected, then the specification is open to spatial elements.

The most commonly used spatial model include the spatial lag model (SLM):

$$y=\rho Wy+\alpha\iota+X\beta+u \quad (1)$$

where the dependent variable is a vector $(n \times 1)$, Wy represents the endogenous spatial interaction effect and ρ is the spatial lag coefficient, ι is a vector $(n \times 1$ of ones that are associated with the constant α, $X$is a matrix of explanatory variables of the order $(n \times K)$ β is a vector of unknown parameters of the order $(K \times 1)$ and $u=(u_1, u_2, ..., u_n)'$ is a vector of random terms of dimension $(n \times 1)$, where u_i is independent and identically distributed throughout i with a zero mean and σ^2_i. variance.

Another standard spatial model is the spatial error model (SEM):

$$y=\alpha\iota+X\beta+u, u=\lambda Wu+\varepsilon \quad (2)$$

where Wu is the effect of spatial interaction on the error term or the residual spatial effect and the parameter λ is the spatial autocorrelation coefficient for the errors.

The combination of these two models (SLM+SEM) gives rise to the spatial autoregressive with autoregressive error (SARAR) model. Other more complex models can be estimated by including spatial lags in the explanatory variables (WX), an approach known as Durbin models, or by streamlining the cross-sectional model into a spatiotemporal one. For more on these alternatives, see Herrera (2017).

The most widely used model includes the endogenous spatial lag. At least three factors underlie its popularity. First, using an economic scheme from game theory with quadratic payoff functions, the SLM can be justified as an approximation to the simultaneous reaction function between *n* individuals (Brueckner, 2003). Another significant reason is econometric theory, as the inclusion of an endogenous spatial term prevents estimator inconsistency, but if the spatial lag is omitted from errors, all that is lost is efficiency. Finally, there is an empirical reason for estimating the endogenous interaction coefficient: the estimated model yields an estimation of the spatial *spillover* that is very attractive in applied terms.

There is no need for a huge database to estimate these spatial effects. In the following section, I give an example of an estimation of these models using a standard database.

TRADE AND SPATIAL EFFECTS

The gravity model for trade allows us to predict bilateral flows between countries and has been widely used in empirical studies in this area (Feenstra, 2004; Helpman, Melitz and Rubinstein, 2008; Krisztin and Fischer, 2015). I will use a simple specification for the model to highlight how the interdependent nature of trade flows needs to be considered differently to the standard treatment.

The data I use here comes from the study by Martin, Mayer, and Thoenig (2008), in which the gravity model was used to measure how trade flows were impacted by a set of components which were divided into nonpolitical variables (adjacency, distance between the two, language similarity, and colonial ties) and political variables (trade agreements and the existence of a communist regime). The authors used a huge international database of trade between 250 countries from 1950 to 2000. In this exercise, I selected data for the last five years included in the original study (1996–2000) for 16 countries in the Americas.[1] There was no variation in the political variables for this subset, so only the nonpolitical variables were used. The nonspatial gravity model to be estimated is as follows:

TABLE 1
ALTERNATIVE ESTIMATES TO THE GRAVITY MODEL

MODELS	MCO	SLM	SEM	SARAR
ln (origin GDP)	1.208 ***	0.685 ***	1.064 ***	0.686 ***
ln (destination GDP)	0.633 ***	0.609 ***	0.638 ***	0.619 ***
Adjacent	10.733 ***	8.753 ***	7.925 ***	8.732 ***
Adjacencyxln(dist)	-1.052 ***	-0.804 ***	-0.726 ***	-0.797 ***
Language index	0.146	-0.043	-0.286	-0.102
W x ln T	0.365 ***	0.350 ***		
W x u	0.476 ***	0.089		
Constant	-17.354 ***	-12.605 ***	-15.602 ***	-12.663 ***
Origin/destination-control	yes	yes	yes	yes
N	1.280	1,280	1,280	1,280
corr² (ln T, $\widehat{\ln T}$)	82.374	85.452	81.631	85.396
Moran´s 1	13.452 ***			
L M robust error	25.625 ***			
L M robust lag	17.442 ***			

Note: *** denotes significance at 1%. Spatial estimation using G2SLS with robust errors.
Source: Compiled by the author.

$$\ln T_{ijt}=\beta_1 \ln GDP_{it}+\beta_2 \ln GDP_{it}+\beta_3 cont_{ij} +\beta_4 cont_{ij}\times \ln dist_{ij}+\beta_5 leng_{ij}+u_{ijt}$$

where GDP_{it}, GDP_{jt} are the gross domestic products of the country of origin *i* and the destination countryj; $cont_{ij}$ is a dummy variable that takes the value of 1 if the countries border one another and 0 if they do not, $dist_{ij}$ is the bilateral distance between capitals, and $leng_{ij}$ as an indicator for language similarity.

In the authors' original specification, the adjacency criterion was used as a variable that allowed them to capture the effect of geographic proximity and thus control for spatial effects that might be present in the database.

In this exercise, I focused on the control variables that seek to capture the effect of geographic proximity, $cont_{ij}$ and $\ln dist_{ij}$. My specification includes the logarithm of the distance between the capital cities of adjacent countries using a term for the interaction between the adjacency criterion and the distance as measured in kilometers. The adjacency coefficient is expected to be positive and to reflect increased trade flows between adjacent countries, although the distance between capital cities should have a negative impact on these flows.

Table 1 shows the estimated models. There are 103 zero values, and 80 of these values are the outcome of the restructuring needed for the spatial estimation, as the same country needs to be taken as the origin and destination of the products. Depending on the aim of the research project, these 80 values might be different to zero as long as flows within each country are captured. The strategy here is to consider this flow as being equal to 0, a common assumption in these types of model, which are also known as origin/destination models (LeSage and Pace, 2009, chapter 8). In short, there are only 23 values without flows between countries (1.8% of the base), which allows us to make a simple estimate for this data.

GEO BIG DATA
THIS PREDICTIVE TECHNIQUE EMPHASIZES INTERCONNECTIVITY AND GEOTAGGING

The W contact matrix was constructed using the criterion of adjacency between destination countries, standardized by row. This formally captures the notion that the flows associated with a destination market improve or worsen depending on the attraction of adjacent destination markets (for more details, see LeSage and Pace, 2009).

The results reveal the difference between a traditional estimation and a spatial estimation. Both recognize the impact of adjacent countries on the specification of the model by incorporating geographical variables. In the basic model (column 1), the specification rests on the assumptions that observations are independent and identically distributed *(IID)* and the spatial nature of the data is approximated through a proxy variable. Using this specification, the different spatial tests detect that the errors do not fulfill the assumption of independence and the model needs to be respecified.

In the spatial models (columns 2, 3, and 4), spatial dependence between the pieces of data is recognized through the inclusion of the adjacency matrix, in addition to geographical variables. Spatial elements can be seen to be significant in all the models, so their inclusion is relevant. The SARAR model provides more information and demonstrates that the endogenous spatial lag is the effect that needs to be considered. In other words, the OLS model omits a variable that leads to inconsistency in estimates.

Numerous publications have attempted to control spatial interaction as is presented in the OLS estimate but without moving beyond this. The estimates in Table 1 show that this type of strategy is not always useful and greater efforts need to be made in this direction. Furthermore, spatial tests are not usually presented as evidence of effective control in these publications.

FINE TUNING

The growing availability of georeferenced information is an unprecedented opportunity for empirical analysis. However, conventional econometric tools fail to take advantage of this wealth of information. This article discusses a selection of spatial econometric tools that can be used to capture the complexity of such data.

Spatial data has very specific characteristics because it comes from a range of sources. Aggregation alters variability and can affect any inferences that are made. The format of geographic data is also important when different types of spatial data are being combined. Progress in data science will gradually allow us to identify and work with alternative data sources at increasingly disaggregated scales. This will lead to empirical studies like the one presented in this paper becoming ever more detailed, enabling researchers to identify the local actors that export or import goods.

This greater detail will challenge standard statistical analysis techniques, forcing us to acknowledge the peculiarities of this data. This is where spatial econometrics comes into its own, by making it possible to process the information in question. It can be used as a tool regardless of the size of the databases in question, which will probably tend to be of the geo big data type. The empirical example included in this paper demonstrates that classical strategies do not capture the complexity of spatial interactions.

NOTES

[1] The countries included are Argentina, Bolivia, Brazil, Chile, Colombia, Costa Rica, Ecuador, Guatemala, Honduras, Mexico, Panama, Peru, Paraguay, El Salvador, Uruguay, and Venezuela.

REFERENCES

Anselin, L. 1988. *Spatial Econometrics: Methods and Models*, Volume 4. Dordrecht: Kluwer Academic Publishers.

—. 2010. "Thirty Years of Spatial Econometrics." *Papers in Regional Science* 89 (1): 3–25.

Arribas-Bel, D. 2014. "Accidental, Open and Everywhere: Emerging Data Sources for the Understanding of Cities." *Applied Geography* 49: 45–53.

Brueckner, J. 2003. "Strategic Interaction Among Governments: An Overview of Empirical Studies." *International Regional Science Review* 26 (2): 175–188.

Cressie, N. 2015. *Statistics for Spatial Data*. Revised edition. New York: Wiley.

Elhorst, J. 2014. *Spatial Econometrics: From Cross-Sectional Data to Spatial Panels*. Heidelberg: Springer.

Feenstra, R. 2004. *Advanced International Trade: Theory and Evidence*. Princeton: Princeton University Press.

Helpman, E., Melitz, M., and Rubinstein, Y. 2008. "Estimating Trade Flows: Trading Partners and Trading Volumes." *The Quarterly Journal of Economics* 123 (2): 441–487.

Herrera, M. 2017. "Fundamentos de econometría espacial aplicada." *MPRA Paper No. 80871. Munich: University Library of Munich, University of Munich*.

Jean, N., Burke, M., Xie, M., et al. 2016. "Combining Satellite Imagery and Machine Learning to Predict Poverty." *Science* 353 (6301): 790–794.

Krisztin, T., and Fischer, M. 2015. "The Gravity Model for International Trade: Specification and Estimation Issues." *Spatial Economic Analysis* 10 (4): 451–470.

LeSage, J., and Pace, R.K. 2009. *Introduction to Spatial Econometrics*. Boca Raton: Chapman and Hall/CRC Press.

Martin, P., Mayer, T., and Thoenig, M. 2008. "Make Trade Not War?" *Review of Economic Studies* 75 (3): 865–900.

Patino, J., and Duque, J. 2013. "A Review of Regional Science Applications of Satellite Remote Sensing in Urban Settings." *Computers, Environment and Urban Systems* 37: 1–17.

Tobler, W. 1970. "A Computer Movie Simulating Urban Growth in the Detroit Region." *Economic Geog-*

Sci-Fi *Macroeconomics*

Daniel Heymann* and Pablo Mira**
IIEP-BAIRES. University of Buenos Aires/National Scientific and Technical Research Council (CONICET)

NEW CHALLENGES ARE AFFECTING THE FIELD OF ECONOMIC ANALYSIS. PUBLIC AND PRIVATE ORGANIZATIONS WITH LARGE DATA STORAGE CAPACITIES AND VAST AMOUNTS OF INFORMATION AVAILABLE TO THEM IN REAL TIME CAN NOW PREDICT AND PROJECT THE DYNAMICS OF A GIVEN ECONOMY. BUT HOW CAN WE COMPLEMENT THE PREDICTIVE CAPACITY OF ARTIFICIAL SYSTEMS WITH HUMAN LABOR?

Some time ago, Isaac Asimov defined science fiction as "the very relevant branch of literature that deals with human response to changes in the level of science and technology" (Stone, 1980).[1] Asimov was a remarkable individual who combined the talent of a great writer with rigorous scientific training and was one of the first to examine the ethical problems associated with "humanoid" robotics (Asimov, 1942). He also said that "I don't know of any science fiction writer who really attempts to be a prophet. [...] Such authors accomplish their tasks not by being correct in their predictions, necessarily, but merely by hammering home—in story after story—the notion that life is going to be different." Anyone attempting to imagine what a future marked by technological change would look like finds themselves in a similar plight: whether they like it or not, they must imagine a context that differs substantially from those on which their own schemes of interpretation and analysis were built.

Uncertainty about the future comes as standard in the current social context. Although the intensity of this uncertainty varies, it affects both concrete decision-making and the study and generation of hypotheses and predictions. In the event of sudden modifications to conditions and behavior, attitudes tend to divide into those who put this down to a new false positive ("the more things change, the more they stay the same") and those who claim that "things are different this time around." History has witnessed huge technical advances (nuclear energy or space exploration, for example) that have not yet changed most people's day-to-day lives[2] as notably as more unforeseen impacts have (such as information and communication technologies).

There are also reasons for thinking that large-scale transformations are underway which will lead to considerable changes in the world of work and production and may even redefine each of our positions in society. The focus these days is on artificial intelligence (AI) and its prospects and potential implications. In a few short years, the issue has become ubiquitous in the media and academia, as is revealed by statistics on the number of references to the related terms in the international press (see figure 1). In the field of economic research, a simple search for through a standard search site (www.nber.org) for documents that mention

AI returns no fewer than 20 records from the last few months alone.

"AI" may refer to a broad range of things. This goes from relatively simple, focused instruments designed to carry out tasks such as recognizing certain specific patterns (signatures on a document, for example) to complex procedures able to outperform humans at activities that require sophisticated reasoning skills. The future promises beings with general intelligence skills that are far above those of Homo sapiens sapiens.[3] Recent years have been marked by intense progress[4] in the development of AI. According to experts (LeCun, 2018), the race is on to find trained, unmonitored systems that are capable of nonspecialist intelligence and are endowed with some kind of common sense.[5] This comes very close to the aspiration of one day being able to build artificial devices that emulate, and perhaps improve on, economic agents made of flesh, blood, and neurons to carry out tasks that require advanced cognitive skills.

How and when this might occur, and what consequences it might have, is not clear, especially for mere economists. Regardless, these notes present reflections on possible economic effects of AI in the form of an essay and include some brief remarks on our own experiences regarding the potential impacts on our field of study. Given the nature of the exercise, we have allowed ourselves to include some value judgments and occasionally explore conjectural tangents.

USERS AND CONSUMERS

AI is an example of an innovation that is already in widespread use. It seeks to transform production and has implications for labor, distribution, specific sectors of the economy, and geographical regions. It will also bring changes to consumption and social connections (Brynjolfsson and McAfee, 2014). Goods and services that have not existed up till now will certainly emerge, opening up human employment opportunities as a consequence of productivity gains. However, many people's ways of life will be called into question.

FIGURE 1
REFERENCES TO AI IN THE INTERNATIONAL PRESS, 2001–2017
(LOGARITHMIC UNITS)

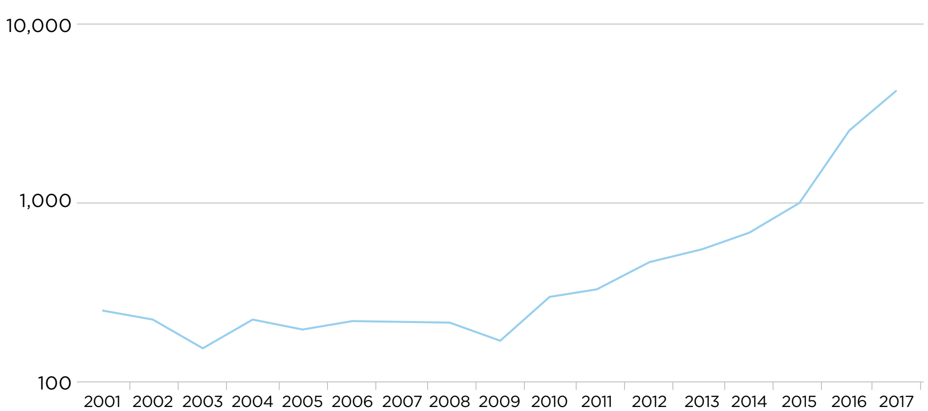

Note: Figure 1 shows the total number of mentions of the terms "artificial intelligence" and "machine learning" per year in articles published in the *Financial Times, New York Times, and Wall Street Journal*.
Source: Compiled by the author based on *ProQuest*.

"Prediction is very difficult, especially about the future," Niels Bohr supposedly once said.[6] Experts' ability to successfully predict future trends and the intrinsic uncertainties in these exercises can be illustrated by examining past attempts to do so. One interesting such example is a 1973 survey in which experts on AI were asked to predict the date when certain technologies would be available on the market (Firschein et al., 1973). A similar exercise was carried out in 2016 (Grace et al., 2017). Table 1 compares these two surveys.

It can be seen that the 1973 responses tended to predict that these developments would spread earlier than they actually have, although there are occasional exceptions. Regardless of the dates by which each achievement was expected to be reached, the changes to the AI agenda are significant. The experts surveyed in 1973 focused on future systems for relatively nonspecific purposes that could be applied to industry and war. Those in the 2016 study highlight future uses for AI in fields such as education and art. Interestingly, the events predicted in 1973 that did not come to pass include the development of AI-based economic forecasting models, which were expected by the early 1990s.

One common area of focus is the effect of technological change on economic growth and the existence and intensity of the effect of decreasing returns from the application of certain production inputs. At present, there is debate around the determining factors for the decline in productivity increases (particularly in the US) despite the growing resources spent on research and development (see, for example, Knott, 2017). At the same time, it has been argued that AI could bring about sustained, rapid economic growth (Yudkowsky, 2013). With regard to its impact on employment, complementarities will certainly emerge between AI and personal skills: this may entail individuals being tasked with assisting algorithms (for example by training or calibrating them) or using their services. However, considerable substitution of the labor force is also to be expected (Frey and Osborne, 2013; see also Acemoglu and Restrepo, 2017a and 2017b; and Korinek and Stiglitz, 2017).

1990
THE YEAR BY WHICH AI-BASED ECONOMIC MODELS WERE PREDICTED TO BE IN PLACE

In a context like the present, assessments of future prospects seems to prompt attitudes that range from over-the-top enthusiasm to apocalyptic outpourings.[7] The optimistic view is that technological change will lead to a state of more or less universal abundance in which AI and humans mutually reinforce one another and in which time spent on backbreaking work will be replaced by creative leisure.[8] Gloomier predictions foresee real threats that a wide variety of human skills will decline, leading to a growing social divide between a privileged few who build on their skills through interactions with AI and reap the profits of this, and a disadvantaged majority whose skills are unsuitable for the new environment.

Conjectures on the potential effects of job displacement have been extremely varied and, to anticipate these effects as a whole, it would be necessary to contemplate the aggregate balance of resources: in practice, artificial workers will need feeding (with energy, mainly) and this will be a real cost in addition to the cost of feeding and maintaining humans, whether they are employed or not. What thus remains to be evaluated is how relative prices would be configured as a consequence and what the incentives for demand of different production factors would be.

Having said this, one widely cited estimate (Manyika et al., 2017) calculated that toward 2030, between 400 and 800 million jobs might be replaced by new technologies and that the jobs performed exclusively by humans would mainly entail physical and emotional interaction with other people, such as social work or care for the elderly (Frey and Osborne, 2013; Grace et al., 2017).

As this is a technology that seems to be headed toward performing tasks that require intensive cognitive and analytical skills, the prospects for job replacement include areas currently occupied by people with high education or skill levels. This poses questions about the specific contents and forms of teaching that will prepare people for productive work in a future where AI is widespread. At the same time, those who stand to benefit most from AI are probably in groups that are already at the top of the wealth pyramid, who can play an active part in generating and making use of innovations in contexts that are prone to the emergence of "winner takes all" effects that then lead to extraordinary profits.

The net impact of AI tends to accentuate economic inequality (IDB, 2017; Korinek and Stiglitz, 2017; Acemoglu and Restrepo, 2017a). When considering collective decisions and policy design, there needs to be focus on the risks these changes pose to those who may lose out from them, especially those who are already in vulnerable situations.[9] There is a very relevant monetary aspect to this: action would have to be taken in the form of public policy if a large share of the population were to find their capacity to generate income seriously restricted. One possible response to a significant contraction in work opportunities is through direct transfers, such as universal income schemes.

An economy in which production is largely performed by machines and programs while large swathes of the population are unemployed, their continued consumption held up by public aid, poses some serious questions that go beyond mere economics. It is worth reflecting whether it is desirable to move toward a social configuration in which large proportions of the population are incapable of contributing to the collective product and become dependent, a situation that may be transmitted from one generation to the next. If this is perceived as too high a social cost, public policies will need to pay close attention to maintaining employability conditions among groups who are vulnerable due to their low income levels and high risk of social exclusion.

On the job supply side, more and better investment in education, with a particular focus on low-income sectors, will be a major component in facilitating access to employment opportunities. Yet even this may not be enough: "virtually every aspect of early human development, from the brain's evolving circuitry to the child's capacity for empathy, is affected by the environment and experiences that are encountered in accumulative fashion, beginning early in the prenatal period and extending throughout the early childhood years" (Shonkoff and Phillips, 2000)[10]. If unfavorable conditions during the early stages of life have persistent effects, it follows that there will be limitations on correcting these if interventions come too late, which translates child poverty into limited future employment capacity. It would therefore be worth trying to sustain sufficient demand for low-skilled workers within the economy, given the intergenerational repercussions that insufficient income levels and social marginality may have on access to skills later in life.

The ubiquitous presence of AI may well be accompanied by considerable asymmetries in the ways in which people relate to it. Different jobs have been identified in the specific sphere of AI use in production, including system trainers; technology communicators, who explain the uses of AI systems to clients; and verifiers, who monitor the performance of AI systems and their compliance with pre-established standards. More generally speaking, social divides may open up between those who have the skills to contribute to building such instruments, those who employ the systems productively, and those who simply use AI services as black box solutions for consumption purposes.

The spread of AI will also affect the international division of labor in various ways. It is probable that significant scale effects will be at work in the design and construction of AI systems. Internet search algorithms, for example, calibrate searches to ensure that "more customers generate more data, which in turn generates more customers" (Goldfarb and Trefler, 2018). Given that

TABLE 1:
PREDICTIONS ON THE AVAILABILITY OF TECHNOLOGIES

GROUP	TECHNOLOGY	**1973** SURVEY	**2016** SURVEY
I	Voice transcription	1992	2022
	Machine translation	1995	2023
II	Universal player/Explain a move in a game	1985	2026
	Robot driver/Driverless truck	2000	2029
	Art valuation system/Generate Top 40 pop song	2003	2029
	Industrial robot/Bipedal robot runner that can beat humans	1980	2030
	Robot household servant/Robot salesperson	2010	2031
	Automatic diagnostician/Robot surgeon	1982	2046
III	Write a high-school essay		2025
	Write a New York Times bestseller		2050
	Mathematical research		2055
	All human tasks		2060
IV	Robot soldier	1995	
	Computer arbiter	1995	
	Insightful economic model	1990	

Note: Years represent the median for all predictions. Group I compares predictions on the same technologies. Group II compares predictions on innovations that are related but not identical. The issues in group III were only included in the 2016 survey, while the ones in group IV were only discussed in 1973.
Source: Compiled by the author based on Firschein et al. (1973) and Grace et al. (2017).

the profitability of having an AI system in a firm depends on the spectrum of applications it works on, and that data that can be used as a learning input for programs can have multiple uses, there will also be economies of scope to new technologies. A third feature of research and development in AI are knowledge externalities, which are associated with direct and immediate access to a range of uncodified but potentially useful ideas and skills. Together, these conditions may favor the geographic concentration of cutting-edge activity, a kind of sequel to what is already happening today in highly specialized hubs in the US and to a growing extent in China and some parts of Europe.

If the big leagues are more or less reserved for heavyweight players, the question remains as to what opportunities are open to peripheral, middle-income economies such as those of Latin America in terms of both possibilities and policy alternatives. In any case, taking refuge in the role of passive consumers of prepackaged innovations is a rather unattractive and even high-risk option because it would imply restricting the capacity to adapt to versatile, unstable environments. Investments in acquiring relevant skills, keeping talent in the country (no small matter given the lure of technological centers), and spreading capacities may facilitate the quest for productive uses of technologies and local contacts and help identify potential areas where local R&D efforts could be successfully applied. This would operate in parallel with the capillary-like spread of these technologies in production applications.

FIGURE 2
GREECE: PROJECTED AND ACTUAL VALUE OF GDP

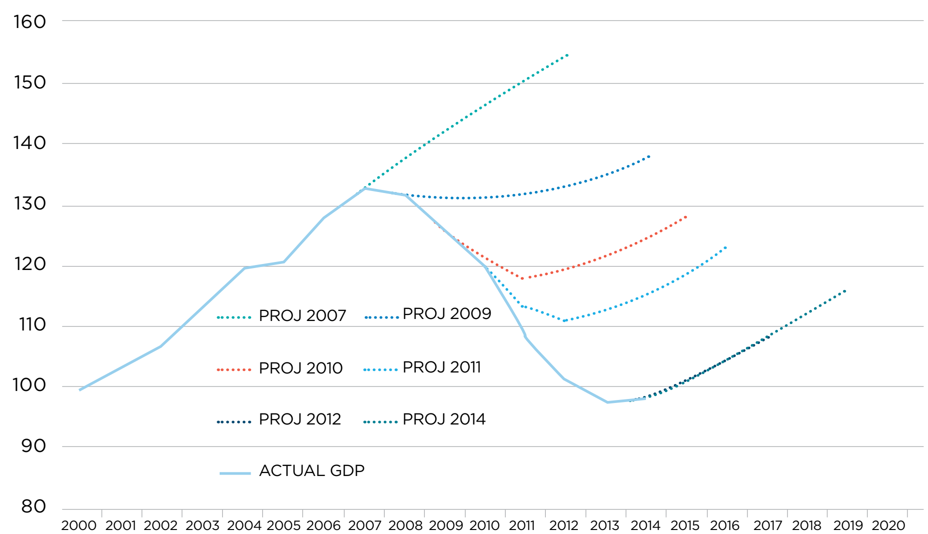

Nota: Each dotted line indicates the future trend that was estimated that year. The unbroken line shows the actual evolution of GDP.
Source: FMI, *World Economic Outlook*.

ECONOMIC DECISIONS

It seems natural for systems designed to process information to be applied to economic decision-making. AI's role in this context may vary considerably, ranging from providing data (the capture of which may be simple or complex) to making predictions and even recommending strategies. In the extreme, AI may be capable of directly implementing actions chosen by a system and becoming a decision-making agent with a certain degree of autonomy.

This has already begun to happen and is being driven by primary stakeholders. There are investment funds that are advertised as being managed by IA (see Kumar, 2017). Loan applications are authorized or rejected every day based on automated credit scoring systems. The giants of the IT world (and other players that are not as huge) exploit vast amounts of data from online searches and transactions to identify behavior patterns that allow them to increase their profits by fine-tuning their services and pricing. Academic analysis has begun to explore the use of AI as an instrument to improve the design of transaction and hiring mechanisms (Milgrom and Tadelis, 2018).

The growing sophistication of the systems themselves and the scope of their application may lead to core aspects of human users' decisions being increasingly delegated to them. In economic contexts, these circumstances naturally give way to agent-principal problems, although with an artificial agent the interaction takes on specific features. In a typical relationship, the principal relates to the agent in terms of their own interests and objectives and has greater information and knowledge in the relevant spheres of interaction. What sets the relationship between humans apart is that the agent has a similar cognitive apparatus to the principal and also has the capacity for introspection. This implies that, on the one hand, the agent can behave strategically and, on the other, that they can account for their actions or recommendations, which are explained by their underlying reasoning. In the current state of AI systems based on neural networks, the relationship could, in principle, include the alignment of incentives if the system were trained for a purpose that is compatible with the desires of the human principal.[11] But advice or recommendations from a system with no capacity for explanation would be very much like that of an oracle, which would affect its perceived reliability. Conversely, if an AI with capacities that far outstripped those of humans for the task at hand had the capacity for introspection, evaluating the outputs generated by the system could lead to a corresponding dilemma: the AI would be able to justify its output but its argument would be beyond the human enquirer's comprehension.

These difficulties may restrict or condition the spread of intelligent contracts with a certain degree of complexity that have been designed using AI. People who receive an offer to take part in arrangements of this sort may be cautious around the information asymmetry that favors the party making the offer because they will not be able to rationalize why they wish to include a given clause. The problem of justification may lead to legal ramifications if one of the parties to a contract drafted by AI files a claim after the fact invoking an arbitrary or discriminatory clause (Bostrom and Yudkowsky, 2014).

Up to this point, this paper has referred to isolated AI systems that operate in a given environment. However, how economies evolve would stem from the collective result of the interactive behavior between people and different types of AI ecosystems that are applied in different areas.

Learning on the part of economic agents induces behaviors that then shape the landscapes in which each agent operates.[12] This is true for both humans and machines. Both, in their way, would also be affected by the problem of incitement and would seek to use past information to project futures that might not necessarily reproduce previous patterns. The emergence of AI as an influential agent or factor in economic decisions has the potential to generate significant changes of this sort. These changes would have particular features if the new agents add greater cognitive capacity than was previously applied. In contexts like this, asymmetries tend to emerge (or become more marked) between the quality of the decisions made by those who have access to artificial systems (and, depending on the context, those who are able to interact with them) and those who are not. At the same time, the changes in the economic context may accelerate due to collective learning and adaptation processes within AI systems, especially if these are more responsive than human agents.[13]

The influence of AI may have significant effects on macroeconomic performance, which hinges on players' perceptions and expectations. In particular, there is the question of how AI might affect the emergence of systemic crises.

As the breaking of contracts and promises is one of the core features of such crises, they are intrinsically associated with widespread frustration of expectations, which manifests itself as a markedly different evolution of incomes as had previously been expected. This is illustrated in figure 2, which shows the contrast between projections and actual aggregate output in Greece, whose economy underwent a severe crisis. As in other similar cases, at a time of reduced income, Greece was forced to face up to debts that it had taken on when reputable estimations suggested a prosperous future.

The development of AI systems may lead to individual decision-making processes becoming increasingly sophisticated and well-informed. However, at the same time, they have the potential to increase the complexity of the environment in which these systems themselves act. There is no guarantee that collective errors that have macroeconomic implications will not appear along the way, as has happened on multiple opportunities with human decision makers, especially during times of economic and technological transition. In other words, AI is unlikely to make crises a thing of the past.

The rise of AI will undoubtedly have repercussions for research and analysis, and economics will not be exempt from this: the field will see new relationships of complementarity and substitution between human work and that performed by artificial systems, and there will probably be a noticeable shift in the way analyses are conducted.

If AI is ever capable of perceiving irony, it may smile at a hypothesis that is commonplace within current economic analysis, namely that flesh-and-blood economic agents are strict optimizers when they have maximum knowledge of their surroundings. It could be surmised that the analytical schemes that underlie the macroeconomic outputs of AI will tend to differ from general equilibrium models in which agents' behaviors are postulated based on optimization problems. Instead, they will move toward models entailing multiple agents (see Janssen and Ostrom, 2006; Heymann, Perazzo, and Zimmermann, 2013) that interact in fluid environments. In such contexts, it is difficult to formulate procedures for identifying strictly maximizing behaviors. Likewise, the rules for making decisions in these environments derive from the observation of many examples of actual behavior that come from the traces agents leave through their online actions. In this case, what may happen is that the comparative advantage in building and operating models may shift from the restricted sphere of academia to private entities (and public ones) with the capacity for obtaining vast quantities of processable data and information. This may be of particular interest when developing instruments to describe and project macroeconomic developments.

NOTES

* School of Economic Sciences, University of Buenos Aires, Argentina. National Scientific and Technical Research Council (CONICET). Interdisciplinary Institute of Political Economy of Buenos Aires (IIEP-BAIRES), Argentina.

** We would like to thank Santiago Chelala, Daniel Aromí, Santiago Cesteros, Gustavo Montero, and Javier Legris for their comments and contributions. The usual disclaimer applies.

[1] Emphasis in the quoted source.

[2] Reserving judgment on the perhaps not-so-distant future (see National Geographic, 2018; Tollefson, 2018). Within the field of artificial intelligence itself, there have been fluctuations in the speed and direction of progress. For example, interest in neural networks went through repeated slumps before accelerating with modern deep learning techniques.

[3] For these ends, a useful (although necessarily hazy) definition of intelligence might be "that sort of smartish stuff coming out of brains, which can play chess, and price bonds, and persuade people to buy bonds, and invent guns, and figure out gravity by looking at wandering lights in the sky; and which, if a machine intelligence had it in large quantities, might let it invent molecular nanotechnology; and so on" (Yudkowsky, 2013). For a critical evaluation of the notion of intelligence as a scalar variable and the argument according to which machines will soon "outdo human intelligence," see Kelly (2016).

[4] Some of these advances have had considerable public repercussions, such as when a robot passed the written test on the entrance exam people must pass to be able to practice medicine in China (Si and Yu, 2017) or when the program AlphaGoZero beat human players and other programs after training without using human data (see https://deepmind.com/research/alphago/). There have also been warnings on the limitations of algorithms in their current state (LeCun, 2018; Hofstadter, 2018).

[5] LeCun (2018) defines "common sense" as the ability to draw on past experiences and make use of information that is not formally structured to fill in the blanks in descriptions of complex environments and in our knowledge of these and use this to make decisions.

[6] It may be worth mentioning another of Bohr's memorable remarks: "Never express yourself more clearly than you are able to think."

[7] This is not the first time such polarized positions have been taken—see, for example, Umberto Eco's (1965) discussion of different positions on the effects of mass media.

[8] It may be worth reminding readers of Keynes's (1930) prediction on shorter working weeks for future generations.

[9] This text does not discuss the implications for tax policies that derive from the possible concentration in ownership of intelligent machines. For more on this, see, for example, Korinek and Stiglitz (2017) and Guerreiro, Rebelo, and Teles (2017).

[10] Cited in Heckman (2006), who adds: "cognitive, linguistic, social, and emotional competencies are interdependent; all are shaped powerfully by the experiences of the developing child [...] Human abilities are formed in a predictable sequence of sensitive periods, during which the development of specific neural circuits and the behaviors they mediate are most plastic."

[11] It is beyond the scope of this paper to consider the dilemmas that will arise around AI systems that can select their own objectives. Fears about this possibility have been expressed by prominent figures such as Hawking (Cellan-Jones, 2014) but dismissed by others (Yudkowsky, 2008).

[12] Lucas's (1976) well-known critique is an example of this proposition, with the (unnecessary) addition of an ill-defined hypothesis on the rationality of expectations (see Heymann and Pascuini, 2017).

[13] Of course, AI could and should learn about the fallibility of simple extrapolations under fluid conditions and the aggregate effect of such learning. It is hard

to predict what properties a system driven by this type of development would have.

REFERENCES

Acemoglu, D., and Restrepo, P. 2017a. "Low-Skill and High-Skill Automation." NBER Working Paper No. 24119. Cambridge: National Bureau of Economic Research.

—. 2017b. "Robots and Jobs: Evidence from US Labor Markets." NBER Working Paper No. 23285. Cambridge: National Bureau of Economic Research.

Asimov, I. 1942. "Runaround." In: *I Robot*. New York: Gnome Press. (1950).

Bostrom, N., and Yudkowsky, E. 2014. "The Ethics of Artificial Intelligence." In: K. Frankish and W.M. Ramsey, editors. *The Cambridge Handbook of Artificial Intelligence*. Cambridge: Cambridge University Press.

Brynjolfsson, E., and McAfee, A. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W.W. Norton & Company.

Cellan-Jones, R. 2014. "Stephen Hawking Warns Artificial Intelligence Could End Mankind." BBC News. December 2.

Eco, U. 1965 (2000). *Apocalypse Postponed.* Bloomington: Indiana University Press.

Firschein, O., Fischler, M.A., Coles, S.L., et al. 1973. "Forecasting and Assessing the Impact of Artificial Intelligence on Society." *Proceedings of the Third International Joint Conference on Artificial Intelligence (IJCAI-73)*. California: Stanford University.

Frey, C.B., and Osborne, M. 2013. "The Future of Employment. How Susceptible Are Jobs to Computerisation." Working Paper. Oxford: Oxford Martin School and University of Oxford.

Goldfarb, A., and Trefler, D. 2018. "AI and International Trade." NBER Working Paper No. 24254. Cambridge: National Bureau of Economic Research.

Grace, K., Salvatier, J., Dafoe, A. et al. 2017. "When Will AI Exceed Human Performance? Evidence from AI Experts." *arXiv:1705.08807*.

Guerreiro, J., Rebelo, S., and Teles, P. 2017. "Should Robots be Taxed?" NBER Working Paper No. 23806. Cambridge: National Bureau of Economic Research.

Heckman, J.J. 2006. "Skill Formation and the Economics of Investing in Disadvantaged Children." *Science* 312 (5782): 1900–1902.

Heymann, C.D., and Pascuini, P. 2017. "On the (In) consistency of RE Modeling." *Anales. Asociación Argentina de Economía Política. 52nd Annual Meeting.* Buenos Aires: Asociación Argentina de Economía Política

Heymann. D., Perazzo, R., and Zimmermann, M. 2013. *Economía de fronteras abiertas. Exploraciones en sistemas sociales complejos*. Buenos Aires: Teseo-UdeSA.

Hofstadter, D. 2018. "The Shallowness of Google Translate." *The Atlantic*. January 30.

IDB. 2017. *Robotlution The Future of Work in Latin American Integration 4.0.* Buenos Aires: IDB/INTAL.

Janssen, M.A., and Ostrom, E. 2006. "Empirically Based, Agent-Based Models." *Ecology and Society* 11 (2): 37.

Kelly, K. 2016. *The Inevitable: Understanding the 12 Technological Forces That Will Shape Our Future*. New York: Penguin Books.

Keynes J.M. 1930 (2010). "Economic Possibilities for Our Grandchildren." In: *Essays in Persuasion*. London: Palgrave Macmillan.

Knott, A.M. 2017. "Is R&D Getting Harder, or Are Companies Just Getting Worse at It?" *Harvard Business Review.* March 21.

Korinek, A., and Stiglitz, J.E. 2017. "Artificial Intelligence and Its Implications for Income Distribution and Unemployment." NBER Working Paper No. 24174. Cambridge: National Bureau of Economic Research.

Kumar, N. 2017. "How AI Will Invade Every Corner of Wall Street." *Bloomberg*. December 5.

LeCun, Y. 2018. "Deep Learning and the Future of Artificial Intelligence." Green Family Lecture Series, Institute for Pure and Applied Mathematics, UCLA. Los Angeles, California. February 5.

Lucas, R.E. 1976. "Econometric Policy Evaluation: A Critique." *Carnegie-Rochester Conference Series on Public Policy*. 1: 19–46.

Manyika, J., Lund, S., Chui, M., et al. 2017. *Jobs Lost, Jobs Gained: Workforce Transitions in a Time of Automation*. McKinsey Global Institute.

Milgrom, P.R., and Tadelis, S. 2018. "How Artificial Intelligence and Machine Learning Can Impact Market Design." NBER Working Paper No. 24282. Cambridge: National Bureau of Economic Research.

National Geographic. 2018. "Stephen Hawking on Space Exploration." *Live from Space*.

Shonkoff, J., and Phillips, D., editors. 2000. *From Neurons to Neighborhoods: The Science of Early Child Development*. Washington, DC: National Academies Press.

Si, M., and Yu, C. 2017. "Chinese Robot Becomes World's First Machine to Pass Medical Exam." *China Daily.* November 10.

Stone, P. 1980. "Science, Technology and Space: The Isaac Asimov Interview." *Mother Earth News.* September/October.

Tollefson, J. 2018. "MIT Launches Multimillion-Dollar Collaboration to Develop Fusion Energy." *Nature* 555 (7696): 294–295.

Wilson, H.J., Alter, A., and Shukla, P. 2016. "Companies Are Reimagining Business Processes with Algorithms." *Harvard Business Review.* February 8.

Yudkowsky, E. 2008. "Artificial Intelligence as a Positive and Negative Factor in Global Risk." In: N. Bostrom and M.M. Ćirković editors. *Global Catastrophic Risks*. New York: Oxford University Press.

—. 2013. "Intelligence Explosion Microeconomics." Technical Report No. 2013-1. Berkeley: Machine Intelligence Research Institute.

Exponential Trade

ARTIFICIAL INTELLIGENCE IS REVOLUTIONIZING THE WAY SUPPLIERS AND BUYERS CONNECT WITH ONE ANOTHER IN RETAIL TRADE, THE AUTOMOTIVE INDUSTRY, THE PRIMARY SECTOR, THE FINANCIAL SECTOR, AND THE WORLD OF LOGISTICS. HOW TO MAKE THE PRODUCTIVITY LEAP INTO THE NEW DIGITAL AGE.

E-commerce in in the Age of Bots

Carlos Chesñevar and Elsa Estevez
National University of the South and CONICET

WITH THE SPREAD OF BOTS, E-COMMERCE IS ENTERING A NEW ERA. THIS ARTICLE EXAMINES HOW BOTS CAN BE USED TO FACILITATE PROCESSES FOR ONLINE SHOPPING, TARGETED MARKETING FOR ONLINE SALES, AND SUPPORT SERVICES FOR CUSTOMERS AND POTENTIAL BUYERS AS PART OF THE SALES AND AFTER-SALES PROCESS.

Artificial intelligence (AI) is a cross-cutting discipline that includes many areas and impacts complex systems as varied as autonomous vehicles, recommender systems, intelligent decision-making, and online searches, among others. There are currently many applications of AI in different areas, including devices that help us go about our daily lives, such as personal assistants like Alexa[1] and Siri[2]; intelligent medical apps that allow cancer to be diagnosed and detected ahead of time[3]; smart systems to handle simple legal matters, such as Compas[4] and Prometea[5]; and driverless vehicles that use AI to receive and interpret contextual information and imitate the way humans drive.

E-commerce is an area of particular interest for applications of AI, one that stands to benefit from applications that would help sellers better understand customer needs and target their sales efforts based on consumer preferences. At the same time, the potential risks associated with the use of AI in e-commerce are limited: the greatest potential problems do not pose threats to human life but simply entail the possibility of making an inappropriate recommendation to a customer or selling them an unsuitable product and so on.

The increased availability of online payment and sales platforms are helping e-commerce operators use AI techniques to profile consumers and apply this information to improving and increasing the potential of different business models. AI has benefited different online business models through specific algorithms that aim to draw intelligent conclusions that enable companies to stay ahead of public demand by providing solutions that pre-empt customers' specific needs.

A 2016 survey[6] of US marketing managers or professionals in similar positions in firms with over 250 employees found that most believe AI has enormous potential for transforming sales and marketing although they are also wary of how these techniques will be introduced into their operations. About 80% of them believed that AI would revolutionize marketing in the next five years but only 26% felt they understood how AI is used and only 10% were already using it. The same survey identified the main challenges and benefits of AI applications (figure 1).

In the next section, we explain the main AI techniques used in machine learning and then look at two particular ways in which e-commerce could be bolstered by uses of AI: targeted marketing for online sales and support services for customers and potential buyers as part of the sales and after-sales processes. Finally, we look at the prospects for AI in e-commerce and sum

A recent study by Ovanessoff and Plastino (2017) has the potential to increase economic growth in South America by one percentage point of gross value-added. The study analyzed the impact of AI use in five economies—Argentina, Brazil, Chile, Colombia, and Peru—which together account for 85% of economic production in South America. Two scenarios were analyzed, one based on current expectations for annual growth and another considering the economic growth expected after the impact of AI has been absorbed by the economy. The two scenarios were calculated for 2035 and it was observed that in the second scenario, the five countries' economies had grown by US$674 billion, or approximately 13%, with Chile and Peru benefiting the most (15%) followed by Brazil and Colombia (13%) and Argentina (9%).

MACHINE LEARNING

One of the major areas in which AI has developed over the last decade is so-called machine learning, which refers to a set of computational techniques that allow complex predictive models to be built from large datasets. Machine learning is the computational base used for data mining (processes that attempt to find patterns in large volumes of information) and business intelligence, and the two use different techniques. Both use databases containing known information which is used to automatically build a predictive model to classify and create associations between new pieces of information. Some of the more widely used machine learning techniques include:

Neural networks: these imitate the architecture of human neurons and the connections between them to reach conclusions, which are subject to a certain level of probability. These networks allow us to solve highly complex problems such as recognizing a human face in a photograph, but they also require high levels of training using known information. This training enables the values of each neuron to be calibrated so that the group can work together to make highly reliable predictions.

Decision trees: these can be built automatically through specific algorithms and are used to represent and categorize a series of conditions that happen one after another, in order to solve a problem.

Classification rules: these allow the category or class of a new individual to be determined based on different known characteristics. For example, a rule that determines whether a customer could potentially be approved credit or a loan based on their monthly salary, previous loans, employment situation, and so on.

Association rules: these allow patterns or regularities to be discovered in a database by connecting one group of characteristics with another. The typical example of this is so-called market basket analysis, in which the average consumer's shopping pattern is automatically inferred by analyzing thousands of shopping baskets—for example, someone who buys bread and milk typically buys butter and jam, too.

The huge quantities of digital information that are available online led to the coining of terms such as "big data" to refer to datasets so large that traditional computing applications need to use special AI algorithms (based on the techniques described above) to find recurring patterns within this data (Marr, 2016). Based on this data and depending on the field of application and type of problem to be solved, other techniques may also be used in addition to the ones described above, or these may be used in combination.

TARGETED MARKETING

Recommender systems are one of the most widespread applications of AI in the world of e-commerce. These are data filtering systems that work based on different types of information (films, music, books, news, images, product descriptions, etc.) that a particular user is interested in. A recommender system generally compares the user's profile with baseline characteristics for the items in question and seeks to predict how a user would rank a given item that the system has not yet considered (Scholz et al., 2017). These characteristics may be based on information about the user's relationship or proximity to the issue or their social environment.

Recommender systems enable automatic customization of online shopping sites, which increases sales and potential customer numbers by transforming visitors into new consumers. It also groups similar products that may be of interest to the consumer and thus increases their brand loyalty. It also consolidates customer loyalty by demonstrating that the company is aware of consumers' different needs and preferences and is trying to help them by guiding their searches towards the products that interest them most. As a consequence, the consumer will end up going back to the same online shopping site after using a competing site because the system has "learned" their tastes and needs without their having to explicitly state them, which naturally increases customer satisfaction.

What factors are used to make AI-based recommendations? Users' demographic data plays a key part in recommender systems—these attributes affect recommendations and may include the customer's age group, their gender, hobbies, the people they know, and so on. A potential consumer's preferences are measured in different ways, such as by counting the likes or ratings that they give a certain product. The amount of time a visitor spends on a

FIGURE 1
CHALLENGES AND BENEFITS OF AI APPLICATIONS IN SALES

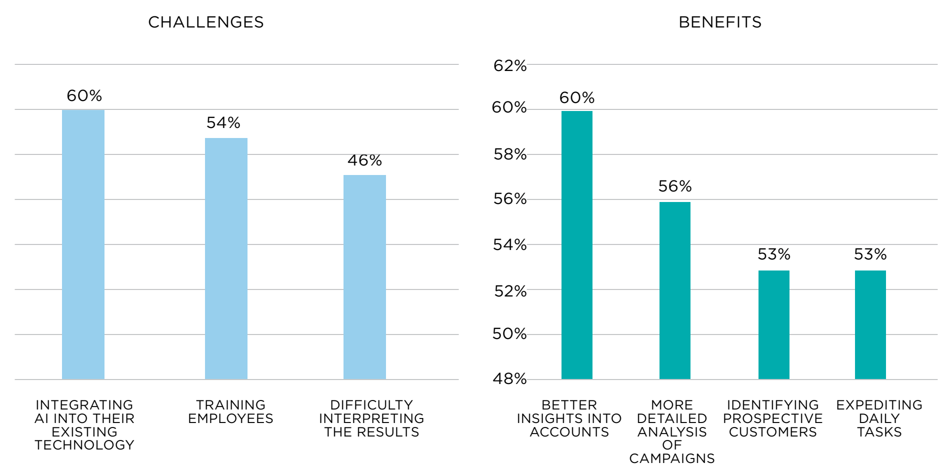

Source: Compiled by the authors based on data from the DemandBase survey (2016).

webpage reading about a product is also an implicit indicator of their interest as a potential buyer. Although these implicit indicators may be hard to obtain, they often provide information that the user may not be willing to give up on their own account but which can be captured without direct action or explicit approval.

Successful examples of the power of recommender systems include sites like Amazon and Netflix. Amazon (Smith and Linden, 2017) is perhaps the most representative case. The company was started two decades ago with the aim of offering every user a personalized bookshop that is tailored to their needs. Every visitor to Amazon's site sees it differently depending on their interests. From a catalogue of hundreds of millions of items, Amazon's recommendation system chooses a small number that are of potential interest to the user based on their current searches and past behavior and those of other users who have bought similar products. The AI algorithms underlying recommender systems help build models in which customers implicitly and anonymously help one another to create the best possible product selection. As Smith and Linden (2017) point out, Amazon launched its collaborative recommendation filtering system as early as 1998, bringing about a true revolution in the automation of recommendations on an unforeseen scale and encompassing millions of customers and products. The Amazon algorithm spread into many other online platforms including YouTube and Netflix. The resulting success was partly due to the scalability and simplicity of the algorithm and the fact that users could easily understand why they were being recommended a particular product.

Netflix is a more recent but equally interesting case. Over 80% of the content that people watch on Netflix is suggested by its recommender system. This means, implicitly, that most of what an average Netflix user chooses to watch is the outcome of an AI-based black box algorithm. As with Amazon, Netflix users machine learning algorithms to discover consumers' tastes and make recommendations. In an interview with *Wired* magazine (Plummer, 2017), Todd Yellin, vice president of product innovation at Netflix, said that the company's underlying business model can be thought of as "a three-legged stool: 'the three legs of this stool would be Netflix members; taggers who understand everything about the content; and our machine learning algorithms that take all of the data and put things together.'" As Yellin explains in the interview, Netflix handles "around 250 million active user profiles. 'What we see from those profiles is the following kinds of data—what people watch, what they watch after, what they watch before, what they watched a year ago, what they've watched recently and what time of day.' This data forms the first leg of the metaphorical stool. This information is then combined with more data [...] gathered from dozens of in-house and freelance staff who watch every minute or every show on Netflix and tag it. The tags they use range massively from how cerebral the piece is, to whether it has an ensemble cast, is set in space, or stars a corrupt cop. We take all of these tags and the user behavior data and then we use very sophisticated machine learning algorithms that figure out what's most important—what should we weigh,' Yellin says. [...] 'What those three things create for us is "taste communities" around the world. It's about people who watch the same kind of things that you watch.'"

TABLE 1
GROWTH SCENARIOS FOR SELECTED COUNTRIES IN LATIN AMERICA

SCENARIO 1		SCENARIO 2					
	BASELINE GDP	POTENTIAL FOR INTELLIGENT AUTOMATION	INCREASE	TOTAL FACTORY PRODUCTIVITY	TOTAL	DIFFERENCE	% DIFFERENCE
BRAZIL	3452	166	192	74	3884	432	13%
COLOMBIA	587	24	42	12	665	78	13%
CHILE	420	21	30	12	483	63	15%
ARGENTINA	643	13	30	16	702	59	9%
PERU	283	7	28	7	325	42	15%
TOTAL	5385				6059	674	13%

Source: Compiled by the author.

CUSTOMER SUPPORT

As we explained above, AI is contributing to the evolution of e-commerce through recommender systems. Another goal is to make the user experience as simple and direct as possible during online shopping. When real-time actions and reactions are used intelligently, they can make a real difference. There has even been talk of the rise of instant commerce, or i-commerce (Zenith, 2017), as a new trend in traditional e-commerce. This form of e-commerce was born, in part, from the need to sell products directly through social media or other online platforms. The aim of i-commerce is to help consumers solve their problem right here, right now, by making the best options available to them and creating the best possible shopping experience for them.

New AI-based voice-controlled assistant services like Amazon Echo[7] and Google Home[8] are still in their early stages but they enable customers to interact with online shopping sites to consult catalogues and make personalized purchases. These new technologies include *chatbots*,[9] which are proving highly effective at providing customer support. These bots are programs that simulate conversations with real people, using AI to generate automatic responses to user questions. To establish such conversations, standard computer programming techniques allow the use of just a handful of phrases, which prevents real, flexible dialogue from taking place between a potential customer and a salesperson (or between a customer and after-sales support services). AI allows programs to simulate human reactions and responses more realistically, such that chatbots can hold fairly logical conversations that customers may find hard to distinguish from real human interactions. By way of example, figure 2 shows the start of a dialogue with the Messenger marketing chatbot tutorial.

Increasing sales while cutting costs is clearly ideal for any business, and chatbots make this possible. According to Raffath (2016), in the US alone, around US$79 billion is spent on customer service, including wage costs. It is estimated that chatbots could potentially replace of up to 29% of customer service staff, which would represent savings of US$23 billion a year.

The many advantages of chatbots include the fact that they provide customer service 24 hours a day without incurring any human resource costs. Their response times are practically immediate, which benefits i-commerce, and they personalize users' shopping experience by tailoring their responses based on customer data and knowledge about their tastes and preferences.

Current technologies mean that AI-based chatbots provide online shopping assistance that draws on historical customer data in real time to enable some of the following: finding out what customers think of their purchases; strengthening communications between customers and firms by sending messages that are tailored to the customer's habits and preferences; asking a customer if they would like to change the address which their purchase will be sent to, finding out what their preferred delivery times are, and whether there are any restrictions on this; notifying customers when a product is out of stock, when it has come back into stock, or when it is expected to; and suggesting alternative products or ones that go well with their favorite purchases.

Firms need to define a specific business strategy if they plan to add chatbots to their platforms. Chatbots can be used for different basic purposes: to provide information, to collect user experiences, and for sales and services, depending on what the consumer is seeking and what the firm is willing to provide. Companies must clearly set out the needs that they are hoping to respond to using chatbots, and this should be aligned with an action plan that includes chatbot monitoring and follow-up.

As was observed in *ChatBots Magazine* (Itsquiz, 2017), bots can represent brands and act on their behalf with consumers. "In 2016, Microsoft created its Bot Framework[10] to support programmers develop bots for their own apps. Facebook strengthens developers with instruments to generate more structured messages, such as descriptions, call-to-actions, images, and URLs. It helps to speed up the relationship between e-commerce and chatbots and expands Facebook's audience" (Itsquiz, 2017).

When digital platforms came into existence, consumers started to receive information about brands as they interacted with a website. The rise of chatbots has improved these communication possibilities by connecting customers with companies intuitively, similarly to the way that personalized customer service does.

A PROSPECTIVE VISION

According to *ChatBots Magazine* (Itsquiz, 2017), in the future, multiple AI techniques will be used to provide efficient technological solutions. *Forbes* (Columbus, 2017) says that 80% of companies are investing in AI. Researchers think that we are far from witnessing a sole, unified form of AI that can be applied in every sphere. We need to take gradual steps that consolidate different approaches (recommender systems, chatbots, intelligent assistants, etc.), all of which will undoubtedly improve in coming years and become more widespread, given the growing adoption of digital media by the general population.

Another major change that is anticipated in the near future is the revolution and adoption of spoken language as a means of communicating with machines. The technologies available through Siri (Apple) and Google Now are relatively recent and the way we use them in everyday life is largely limited to specific questions about the weather, asking for GPS directions, or finding out about the traffic where we live at a certain time of day. In coming years, chatbots are expected to include emotion recognition skills (an area of AI that is only just developing), such that a chatbot can handle a customer request and identify from the person's tone of voice whether they are irritated, bored, or happy to interact with it (and interpret the relationship between this mood and the brand that the chatbot represents).

The rise of chatbots in the world of e-commerce will mean that there are more and more human interactions with machines and that the interfaces used for this kind of communication will be improved. Many brands, such as 7-Eleven[11] and Lego[12], have developed chatbots to contact consumers through their preferred communication channels (such as personal Facebook accounts). These technologies are clearly going to disrupt traditional understandings of e-commerce.

According to the specialist consultancy firm Gartner (2017), by 2020 AI will create more jobs than have been destroyed by its emergence, which represents a real shift in employment dynamics. According to the firm, the number of jobs affected by AI will gradually change in 2018 and 2019, with employment increasing in the health, education, and public sectors and decreasing in manufacturing. A company press release (Gartner, 2017) argued that "AI will improve the productivity of many jobs, eliminating millions of middle- and low-level positions, but also creating millions more" in management and automation. Svetlana Sicular, research vice president at Gartner, stressed that the greatest benefit that could come from AI is AI augmentation, "a combination of human and artificial intelligence, where both complement each other." In this sense, e-commerce is expected to be one of the disciplines that will most benefit from the adoption of these technologies. Likewise, the traditional online shopping model is expected to expand to become technologically more superior, with greater potential for companies and customers to provide highly personalized, high-quality sales services through the use of AI technologies.

In recent years, AI has become a major player in different fields, expanding the potential of existing technologies and creating new application development niches. As we have described in this article, e-commerce is not exempt from these changes. Different AI ap-

FIGURE 2
MESSENGER MARKETING CHATBOT

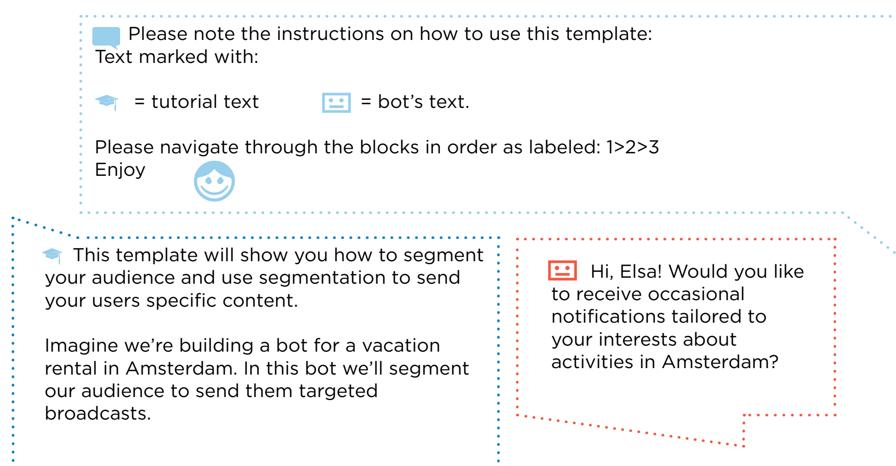

Source: Messenger Chat Bot Marketing, Facebook.

plications, including recommender systems, chatbots, and automatic handling of complaints, are helping to improve and leverage prospects for e-commerce companies. This is taking their products into new market segments and creating new metaphors for building customer brand loyalty. The use of AI in e-commerce is fundamentally based on creating automated interactions between the consumer and the service provider that are increasingly similar to human interactions but that take place in the digital world. At the same time, these developments imply ethical challenges in relation to personal data security and the commercial use of this data within a clearly defined context. In this sense, practices are moving faster than regulatory frameworks can be developed. Despite this, AI is here to stay and is extending current technological capacities in different directions. E-commerce stands to benefit enormously from this spread, which will undoubtedly lead to new business models and other possibilities in the future.

NOTES

[1]Alexa is a voice-based cloud service developed by Amazon that is now included in different Amazon devices.

[2]Siri is a voice recognition service for mobile devices developed by Apple.

[3]The Computer Science and Artificial Intelligence Lab at the Massachusetts Institute of Technology (MIT), the Massachusetts General Hospital, and the Harvard Medical School are researching the use of AI to improve breast cancer detection and diagnosis. The three teams are working together to develop an AI system based on machine learning to predict whether "a high-risk lesion identified on a needle biopsy after a mammogram will upgrade to cancer at surgery" (Conner-Simons, 2017).

[4]Compas is an AI-based risk assessment system used by the Wisconsin Supreme Court (Tarantola, 2017).

[5]Prometea is an AI-based system specifically designed to predict solutions to simple legal cases. The system is used in the Attorney General's Office of the Autonomous City of Buenos Aires in Argentina and at the Inter-American Court of Human Rights. For more information, see the article on this issue in this publication.

[6]DemandBase survey (2016), in association with Wakefield Research.

[7]Amazon Echo is a voice-controlled device whose functions include answering questions and music playback.

[8]Google Home is a voice-controlled device that allows users to interact with services through Google's personal assistant.

[9]"Bot" is short for "robot" and refers to a software program that carries out repetitive tasks online.

[10]Microsoft Bot Framework is a tool for developing bots.

[11]The 7-Eleven chatbot tells customers about special offers and the stores nearest to them.

[12]The Lego chatbot helps users select the most appropriate product to give as a gift.

REFERENCES

Columbus, L. 2017. "80% of Enterprises Are Investing in AI Today." *Forbes*. October 16.

Conner-Simons, A. 2017. "Using Artificial Intelligence to Improve Early Breast Cancer Detection." *MIT News*. October 16.

DemandBase. 2016. "80 Percent of all Marketing Executives Predict Artificial Intelligence Will Revolutionize Marketing by 2020." *DemandBase.* December 13.

Gartner. 2017. "Gartner Says by 2020, Artificial Intelligence Will Create More Jobs Than It Eliminates." Press Release. December 13.

Itsquiz. 2017. "Chatbots for Retail and E-Commerce—Part Three." *ChatBots Magazine*. March 13.

Kabiljo, M., and Ilic, A. 2015. "Recommending Items to More than a Billion People." *Code*. June 2.

Naughton, B. 2016. *Big Data in Practice: How 45 Successful Companies Used Big Data Analytics to Deliver Extraordinary Results*. Chichesterand Malden: John Wiley and Sons Ltd.

Ovanessoff, A., and E. Plastino. 2017. "Como la Inteligencia Artificial puede generar crecimiento en Sudamérica." Accenture.

Plummer, L. 2017. "This Is How Netflix's Top-Secret Recommendation System Works." *Wired.* August 22.

Raffath, H. 2016. "Chatbot Is the Future: 80% Companies Want to Integrate It by 2020." *DazeInfo.* December 15.

Scholz, M., Dorner, V., Schryen, G., et al. 2017. "A Configuration-Based Recommender System for Supporting E-Commerce Decisions." *European Journal of Operational Research* 259 (1): 205–215.

Smith, B., and Linden, G. 2017. "Two Decades of Recommender Systems at Amazon.com." *IEEE Internet Computing* 21 (3): 12–18.

Tarantola, A. 2017. "Watson Is Helping Heal America's Broken Criminal-Sentencing System." *Engadget.* August 25.

Zenith. "La inteligencia artificial en el *e-commerce*: el *iCommerce* cobra protagonismo." *Bloggin Zenith.* May 18.

Karim Lakhani
Harvard University

“Tolerance of failure is the *key* to development”

FROM THE LABORATORY FOR INNOVATION SCIENCE AT HARVARD (LISH), KARIM LAKHANI WORKS WITH NASA AND OTHER US FEDERAL AGENCIES AT THE FOREFRONT OF INNOVATION. LISH ALSO ORGANIZES ONE OF THE HIGHEST-PROFILE PROGRAMMING COMPETITIONS IN THE WORLD. IN THIS CONVERSATION, HE STRESSES THAT LATIN AMERICAN IMAGINATIONS HAVE A MAJOR ROLE TO PLAY IN SOLVING THE REGION'S PROBLEMS AND ARGUES THAT A LACK OF CAPITAL AND DIFFICULTIES IN ACCESSING NEW TECHNOLOGIES NO LONGER NEED RESTRICT THE GROWTH OF EMERGING COUNTRIES.

How will artificial intelligence (AI) transform our economy?

A core part of most of the services we use, like Google, Facebook, or Netflix, is already AI-enabled. Examples of this include the recommendations that Google makes when we are searching for something, Google Translate, or Netflix's recommendations of other movies we might like to watch. Over the last decade, we have seen a transformation in how data is generated and used to make predictions. While the tech sector has been using this approach intensively for the past three to five years, these capabilities are now becoming broadly enabled for everybody. We have always had data, but we were mostly using it for retrospective analysis, to understand what happened in the past. Now we are using it to predict what will happen in the future.

Which sectors do you think will see the greatest impact?

Autonomous driving, where a lot of the major technological problems have already been solved. The questions that arise now are around the business model for these companies and how cities and governments will react to robotic cars on the streets. That will be the biggest challenge.

Are there any other sectors where we can expect to see major transformations?

I would flip the question around and ask which areas people need to make more predictions in. The answer is a lot of areas. Medical imaging, for instance, is a huge potential area for AI because the people that look at medical images return a high proportion of false positives and false negatives. AI could help flag issues and provide a second opinion. More generally, AI could play an important role in supporting any kind of decision which is data-driven. Marketing campaigns are a good example of this: AI could be used to identify which customers to pursue or to predict which

Karim Lakhani is a professor at the Harvard Business School. He specializes in management and technology and is a founder and codirector of the Laboratory for Innovation Science at Harvard. He is a global expert on crowdsourcing innovation. His most noteworthy publications include:

· "Digital Ubiquity: How Connections, Sensors, and Data are Revolutionizing Business." *Harvard Business Review* 92 (11): 90-99. With M. Iansiti. 2014.

· "Using Open Innovation to Identify the Best Ideas." *MIT Sloan Management Review* 55 (1): 41-48. With A. King. 2013.

· "Using the Crowd as an Innovation Partner." *Harvard Business Review* 91 (4): 61-69. With K. Boudreau. 2013.

· "Open Innovation and Organization Design." *Journal of Organization Design* 1 (1): 24-27. 2012.

customers are going to stay or leave. HR is another area: there is a whole domain on how to apply analytics and machine learning to running an organization, managing teams, and defining who should be hired and who should be promoted. In these sectors, AI and analytics will be extremely helpful.

Will there be any differences in emerging economies?

Technology is in many ways open, free, and available to everybody. The constraints are going to be around data and the ability to use it effectively and adapt organizations to it. At the micro-level, the challenges facing companies will center on how to implement this technology. It's a similar story at the municipal level. For instance, the challenge for cities is knowing what to do when a platform like Airbnb or Uber shows up in your marketplace. What would the smart response be? The same thing will happen with AI in connection with municipal workers, lawyers, and doctors. These changes are going to happen, so smart policy-makers should be thinking about the experiments that they should be running to understand how we need to respond on our own. The problem is no longer about having the capital to access technology. For instance, at my research lab, we launch competitions around computational and algorithm development, and one of our best contenders was the forensic police force in Rio de Janeiro. Again, access to capital is no longer a problem, the challenge is going to be access to data, but since data is native to every country, the opportunities are out there. The real challenge will be the imagination and creativity that people develop to use data and how the policy and regulatory environment adjust to that. I think there is a real opportunity for us to decide to enable experiments in cities and countries to develop better policies.

How does this apply to Latin America in particular?

Part of the discussion for Latin America should be, are we going to build our own hubs or should we have smart policies around innovation and competition making sure that our companies and organizations can succeed within existing hubs? How can we encourage our companies to take part in multiple hubs? How do we make sure that data is not locked up in one organization but is widely available? We need to respond through collective action by companies and also through smart regulation. Latin America and the Caribbean have to move from being users of these technologies to becoming producers. We are no longer in a world where we need large factories, industrial plants, or capital equipment. This is in many ways the lowest-cost technology that has ever been available: data is everywhere, almost every citizen is carrying some sort of smart device, and sensors are very cheap. If the constraint is not capital and we have a good labor force, then the question is, what *are* the constraints? For me, it's about imagination, a culture that can encourage risk-taking and tolerate failure or celebrate it in many ways. A sense of community, people who can work together to change the world. Sometimes people in emerging economies are too modest in their ambitions. Why? They don't need to be.

THE MOST PRIZED ASSET WILL BE DATA. EVERY COUNTRY HAS ITS OWN DATA AND THE IMAGINATION TO USE IT

How can Latin America's SMEs and other companies adapt to these changes?

For companies, the challenges and the opportunities around AI and analytics are about changing business models and transforming how we create and capture value. Part of the innovation coming out of Silicon Valley is business model innovation. Policy-makers and governments need to create experimental zones for these companies, give them the latitude they need to be able to move through the learning curve as quickly as possible by applying analytics. From municipal services to healthcare, there are going to be a range wide of sectors where these technologies are going to be applicable. The challenge centers on how you become a test-bed for innovation so that you and your companies benefit from it.

What difficulties might arise when humans and machines coexist in the same working environment?

In the short term, the story is complementary. In the long term, we have some fundamental questions regarding replacement and what humans will do instead. The best analogy I have about this is photography and what it did to the art scene when it was first launched. It used to be the case that in Western European art, still life was the gold standard. But suddenly, photography arrives and then anyone could create a still life through photographs. The art world was in a mess since a machine could take a better picture than what a human could draw. However, in response to photography, modernism and artists such as Picasso emerged. There is a threshold of creativity around this new technology. I think that we have now very good examples of art made with AI, like music. We don't know yet how will AI impact the world of art, but I imagine that now something similar will happen with AI and jobs. There is a scary part and an exciting part.

RESEARCH ECOSYSTEM

- **Accenture Artificial Intelligence** — accenture.com/us-en/artificial-intelligence
- **AI Now Institute at New York University** — ainowinstitute.org
- **Allen Institute for Artificial Intelligence** — allenai.org
- **Machine Intelligence Research Institute** — intelligence.org
- **McKinsey & Company Artificial Intelligence** — mckinsey.com/global-themes/artificial-intelligence
- **Microsoft Artificial Intelligence** — microsoft.com/en-us/research/research-area/artificial-intelligence/
- **Association for the Advancement of Artificial Intelligence** — www.aaai.org
- **Center for Human-Compatible AI** — humancompatible.ai
- **Getting Cognizant Artificial Intelligence Right** — cognizant.com/perspectives/getting-artificial-intelligence-right
- **Partnership on AI to Benefit People and Society** — partnershiponai.org
- **Princeton Center for Information Technology Policy** — citp.princeton.edu/research
- **Natural Language Processing Group** — pln.famaf.unc.edu.ar
- **Deloitte. Artificial Intelligence in the Public Sector** — www2.deloitte.com/au/en/pages/public-sector/topics/ai-public-sector.html
- **European Association for Artificial Intelligence** — eurai.org
- **Facebook AI Research** — research.fb.com/category/facebook-ai-research-fair
- **Stanford University-Artificial Intelligence Laboratory** — ai.stanford.edu
- **Tech Policy LAB** — techpolicylab.org/research-areas/robotics
- **The AI Initiative, Civic Debate on the Governance of AI** — ai-initiative.org
- **Foresight Institute** — foresight.org
- **Future of Life** — futureoflife.org
- **Future Today Institute** — futuretodayinstitute.com
- **The Future Society** — thefuturesociety.org
- **Artificial Intelligence Research Institute** — iiia.csic.es/en
- **Universität Bremen Artificial Intelligence** — ai.uni-bremen.de
- **Google DeepMind** — deepmind.com
- **IBM AI Research** — research.ibm.com/ai
- **Institute of Electrical and Electronics Engineers** — theinstitute.ieee.org/technology-topics/artificial-intelligence
- **Vector Institute for Artificial Intelligence** — vectorinstitute.ai
- **Wadhwani Institute for Artificial Intelligence** — wadhwaniai.org
- **Waterloo Artificial Intelligence Institute** — uwaterloo.ca/artificial-intelligence-institute/

Urban Techno-transportation

Phil Tinn and Michael Lin
MIT Media LAB

DRIVERLESS VEHICLES, ELECTRIC BICYCLES, AND RIDE-SHARING WILL BE THE COMMON DENOMINATORS OF THE TRANSPORTATION OF THE FUTURE. ARTIFICIAL INTELLIGENCE WILL BE USED TO READ TRAFFIC SIGNS AND CITYSCAPES TO OPTIMIZE MOBILITY WITHIN AN INTERCONNECTED TRANSPORTATION NETWORK. THIS ARTICLE DISCUSSES HOW EXISTING INFRASTRUCTURE CAN BE ADAPTED TO THESE NEW WAYS OF MOVING ABOUT.

Over the course of the 20th century, cities have emerged as veritable centers of social, economic, political and cultural activities, with 66% of the world's population expected to live in them by 2050—an additional 2.5 billion people compared to today's urban population of roughly four billion (UN, 2014). Cities compete internally to attract and support human capital in order to boost their economies. In cities, time is a critical resource for both productivity and leisure. Due to the rapid population increase that has, by and large, outpaced the evolution of transportation infrastructure, the planning and design of transportation policies, services, and technologies today all have a real impact on the socioeconomic well-being of cities and their inhabitants.

Vehicle miles of travel (VMT) by cars, a measure often correlated to GDP growth in the past, is no longer widely assumed to be an indicator—much less a driver—of economic progress (Ecola and Wachs, 2012). In the United States, higher per capita GDP can be observed in states with lower per capita daily VMT (figure 1). Similarly, in Germany, per capita GDP is no longer found to be higher in places where people commute longer distances.

Almost to the contrary, according to an upward mobility study conducted by Raj Chetty and Nathaniel Hendren (2015), short commute times as a neighborhood characteristic have been found to contribute to income mobility more than various other factors such as one's parental income, middle-class origin, and social capital (figure 2). If measures of distance traveled are no longer a guaranteed contributor to the improvement in one's economic and cultural life in cities, what value can reinventing the automobile bring, after all? Cars, the form of personal transport that replaced the horse, have been the technology par excellence that supported and, in turn, were supported by suburbanization and sprawl. What will be their intrinsic utility in the age of urbanization, where roads no longer have sufficient space for them to flow efficiently and where city mayors and inhabitants are losing patience with parking and breathing the air pollutants they release?

RIDE-HAILING

The last few years have seen a rapid adoption of app-based taxis (ride-hailing), shared bikes, and increasingly, autonomous vehicles across the world. In 2017 alone, ride-hailing companies such as Lyft, Uber, Didi, and Grab have raised more than US$9.6 billion, and bike-sharing companies such as Ofo and Mobike at least US$1.3 billion (figure 3). In the first half of 2018, Ford officially launched its bike-sharing service GoBike in San Francisco, and Uber acquired its own for US$300 million. According to the Brookings Institute, autonomous vehicle-related ventures and R&D garnered US$80 billion in investment between 2014 and 2017. Fueled by such a large influx of capital for R&D and expansion of new transportation services around the world, how are these new technologies improving lives and turning into sustainable services?

Ride-hailing services have undeniably helped shift people's travel preference away from car ownership, once a symbol of personal economic well-being. Ride-hailing also verified the demand and market for on-demand, door-to-door mobility services, serving essentially as a prototype for driverless cars that transport passengers between places without their involvement in the vehicle operation itself. But ride-hailing suffers from two deficiencies: first, contrary to the messianic vision that ride-hailing services would alleviate traffic jams in cities, it generates more trips and reduces the use of active transportation (e.g., walking, cycling) and public transit (e.g., buses) (Clewlow, 2017). Second, it's constrained by the legacy footprint of the four-wheel automobile. One can easily see why ride-hailing, as convenient as it appears from the consumer perspective, has yet to contribute to improving the flow of traffic on the road. These two constraints, arguably, may also limit the socioeconomic impact of driverless cars, should the industry continue its current path of merely replacing the driver with the computer without challenging the paradigm of the car itself.

Bike-sharing brings multiple benefits, such as a reduction of carbon dioxide emissions from the transportation sector, which, as it causes 1.8 billion tons of these per year in the US (E360, 2017), has now become the number one source of these emissions. It also brings about a reduction in vehicle footprint (compared to the car's) and an increase in the opportunity for nonbike-owners to cycle to work for exercise.: at an increase in the opportunity for non-bike owners to cycle to work and for exercise. Bike-sharing has become widely adopted around the world in the past decade, with increasing permutations such the dockless system first pioneered by Chinese startups, which purports to "disrupt" bike-sharing by offering a more flexible last-mile vehicle not bound by the fixed location of docking stations. With the help of venture capital funding, dockless systems enabled cities to rapidly demonstrate their participation in the active mobility movement. Dockless bike-sharing, in general, still faces three challenges: its contentious use of public space, the challenge of fleet rebalancing, and long-term financial sustainability.

Dockless bike-sharing frequently results in users parking their bikes in random public places. Increasingly, such behaviors are perceived as a "nuisance" by the general public, creating a new phenomenon referred to by some as "bike pollution." If we park the shared bikes freely on the street today in the name of the "sharing economy," what else might be left freely on the street in the future? In addition to challenging the traditional order of the sidewalk and public spaces, dockless systems still retain the unresolved cost problem of fleet rebalancing that already plagues the dock-based systems. Even with the adoption of predictive analytics using AI to anticipate demand across the city, the activity of fleet re-allocation remains a manual one, requiring labor to drive a truck, van, or trailer to physically collect bikes from zones and stations with excess bikes and offload them to places with empty docks.

For all the love for new travel options offered through bike-sharing and ride-hailing, cities are facing difficult transitions toward ensuring long-term collective interests—notably in the protection of public spaces and public transit systems. Amid this seemingly innovative era in the domain of urban transportation, driven by Silicon Valley and the venture capital world and that is awaiting the arrival of the level 5, fully autonomous car as a savior, what might the public and private sector do in the meantime to ensure

concrete progress toward improving urban transportation? Using examples primarily from MIT, here we outline three approaches to illustrate the way the public and private sectors can work together with research institutions to meet this goal: 1) push the limit of active and shared mobility; 2) bring cocreation to AI systems used for the road; and 3) leverage open data to improve infrastructure readiness.

MOVE MORE WITH LESS

The first and foremost approach for alleviating congestion in cities in the context of rapid population increase is to reduce the per-capita unit of mobility in the city—namely by increasing the attractiveness and efficiency of active and shared mobility services (such as cycling, scooters, and car-sharing). Cycling has a critical role in the emerging constellation of travel options in cities, not just due to its low modal cost (including the costs of operation, parking, and crashes) (Litman, 2018), but also due to the questionable improvement for urban traffic that autonomous cars will bring to cities. As Robin Chase (2018), the cofounder of car-sharing enterprise ZipCar, explains, once the car is "freed of the driver, the marginal cost of moving a car will be insignificant... rather than pay for parking, it will be cheaper to keep the vehicle flowing in traffic or return it home to park for free." The likely result: more cars on the road, more congestion, and more pollution. Shared autonomous cars, however, suggests Chase, can help alleviate that danger. At MIT, we further improve the space efficiency of shared autonomous vehicles by reducing their footprint to that of a bike.

The Persuasive Electric Vehicle (PEV) project at the MIT Media Lab exemplifies this approach to footprint reduction by taking advantage of the falling cost of electrification and AI computation and the new behavior of vehicle sharing to boost the viability of cycling (figure 4).

FIGURE 1
PER-CAPITA VMT AND GDP, US STATES, 2008

Source: Duduta (2010).

Autonomy as a method for reducing the cost of labor is already well underway at major ride-hailing services (e.g., the Lyft–nuTonomy partnership). It has yet to be applied to bike-sharing as a technique to automate the rebalancing of vehicle stocks between places that need more bikes and places that need fewer. Once bike-sharing systems can redistribute bikes automatically, they will be able to offer a door-to-door, ride-hailing-like service. This will increase the convenience factor of renting a bike and thereby lead to more people choosing to cycle. In addition, the future operator can expect to fulfill the same amount of travel needs using significantly fewer bikes. Our preliminary study of a bike-sharing fleet's performance, substituting the traditional two-wheel, nonelectric bikes of Boston's bike-sharing program Hubway with electric, self-driving three-wheel bikes, resulted in a 47% reduction in fleet size while achieving a passenger wait time of less than five minutes (pick-up) for 75% of the trips. This improvement in performance is prior to introducing any machine learning–based predictive rebalancing techniques.

An associated benefit of integrating self-driving technology with bike-sharing is the expanded utility of providing door-to-door, micrologistics services, a new function that can generate additional revenue to potentially subsidize the cost of passenger commutes. While there's an ongoing debate as to whether bike-sharing should qualify as public transit and receive government subsidies, what's irrefutable is the low utilization rate of shared bikes in most cities and the potential opportunity to increase its usage and service scenarios. Bikes that are part of the Paris bike-sharing system, Vélib, are used for an estimated average of 5.3 to 6.7 rides per day, whereas those in New York's system are used for 3.8 to 8.3 trips per day, according to reports from the National Association of City Transportation Officials (NACTO) and the Institute for Transportation and Development Policy (Fillin-Yeh, 2017; ITDP, 2013). Bikes in small to medium-sized American cities, on the other hand, average just 0.5 to 2 rides per day. Over a timespan of 24 hours in which trip duration averages 12 minutes for members and 25 minutes for casual users, most systems' bikes are being used to transport passengers less than 10% of the time. What might become of the remaining 90% if shared bikes become more attractive, door-to-door, and multifunctional? If city centers were to become car-less one day, as many progressive city mayors—from Oslo to Chengdu—are striving to achieve, how would businesses transport inventories and deliver goods to their customers? Can the bikes, outside of peak commuting periods, be used to delivery pastries to cafés in the early morning, mail and lunch boxes at midday, groceries in the evening, and supper in the wee hours? Do we really need a four-wheel Uber sedan to deliver McDonald's if a smaller vehicle can do the same job cheaper, more sustainably, and not be stuck in a traffic jam? What if, in the near future, one no longer needs to drive a car to carry groceries because a lightweight robot can carry them for you while you jog to and from the grocery store to combine exercise and shopping?

TRUST AND COCREATION

The fatal crashes of the Tesla on its autopilot trial and the Uber on its routine autonomous testing in March 2018 injected a dose of much-needed realism into the discourse on adopting AI systems in our everyday transportation systems. The incidents, in other words, introduced some necessary critical distance between the public, the media, and exuberant support for an AI-led, technological utopianism that has been little examined or reflected on up to now. While fatal accidents are neither new nor avoidable for any transportation technology, the Uber and Tesla accidents may be easily perceived by the public as the direct opposite of the promise of the "safer roads" much trumpeted by the industry and the media. Enter the brave old world of insurance, where policies have traditionally been structured to attribute liability to human operators since they account for 90% of motor vehicle accidents. If Uber's self-driving car had no human operator in the driver seat, how would the responsibility for failure be traced under the structure used for traditional motor insurance policies? Was it the Lidar that did not detect the crossing cyclist and inform the operating system in time, or was it a slight delay somewhere in the data transfer hub? Was it a glitch of the code somewhere or even overheating in some part of the system? Or was it a combination of several factors? Without the human operator, how will insurance policies be structured? Will car companies be prepared to insure their own vehicles similarly to how consumer electronics are insured?

The recent Uber accident in Arizona will not be the last time these questions get asked. From the point of view of the general public, the accident also revealed the overall lack of transparency and engagement in the development of a technology that's meant to be used daily, with the family, and on the roads of the neighborhood we live in—that is, in your and my own backyards. How will people come to trust an autonomous vehicle? What are the major players in autonomous vehicle development doing to gain the trust of their users and people in the community?

The PEV project at MIT Media Lab champions the approach of cocreation as a process to enable a greater number of real-life actors—including pedestrians, cyclists, and motorcyclists—to be exposed to, interact with, and give feedback on the autonomous vehicle while it is being prototyped. As early as late 2015, we began deploying the PEV prototype in uncontrolled bike lanes and sidewalks around MIT as a test for initial reception from the student community. By August 2017, we began deploying the prototype in bike lanes and sidewalks that are actively occupied by cyclists and pedestrians to allow potential scenarios in which uninformed persons could encounter the PEV in motion. This provided us the opportunity to have conversations with people in-situ about the technology and about their expectations, fears, and hopes.

To ensure that communities in other parts of the world (that is, in other cultural and urban contexts outside of the MIT ecosystem in Cambridge, Massachusetts) can also develop trust in the autonomous vehicle, we began to conduct pilot tests internationally through our network of collaborating cities and states, which consist of Andorra, Hamburg, Helsinki, Shanghai, and Taipei. The most recent deployment took place in March 2018, at Taipei's Daan Forest Park, where we tested last-mile passenger pick-up scenarios on bike lanes during commuting time, in broad daylight, and at night. We observed the reactions and actions of a large sample of cyclists, pedestrians, and joggers as they came in close contact with the vehicle. These deployments in real-life situations outside of MIT allowed us to confirm several design choices made for the PEV, specifically in its scale, speed, agility, form factor, and resemblance to a familiar bike or baby carriage.

In addition to the question of how the public perceives the vehicle's basic physical features and capability, if AI is to gain people's trust, then it will need to develop at least some basic level of social awareness. By doing so, according to the director of the City Science Group, Kent Larson, the PEV will be able to "understand, predict, and respond to the actions of pedestrians and other road users, and communicate its intentions to humans in a natural and nonthreatening way." At a broad level, the design and implementation of the human-machine interaction (HMI) of a lightweight autonomous vehicle such as the PEV differs from that of autonomous cars primarily in the need to interact with people—mainly pedestrians and cyclists—who are often in close contact and approaching from all different directions outside of the vehicle. The deployment at the Taipei park highlighted several critical scenarios in which some form of socially acceptable communication method would greatly facilitate the exchange between the PEV and people: when the PEV and a pedestrian or cyclist are on a path to potentially collide with one another, how might the PEV signal its awareness and intent to the approaching party? If there are multiple parties approaching, how might the PEV behave differently than it would to a single party? In addition to yielding to pedestrians and cyclists, the PEV's interactions with cars will also require development and testing, given that traditionally motorists communicate their intentions through hand gestures, flashing headlights, and honking. Furthermore, this will need to be studied not just in one social context but in multiple ones across the world.

How an AI system behaves in different social, cultural, and geographical contexts points to the need for further research in effective ways to incorporate social intelligence and ethics into AI machines. Making AI sociable and likable in the eyes of the public requires researchers to go beyond the bounds of their laboratories. Public engagement and cocreation, online or offline, will increasingly become a useful avenue for AI developers.

In addition to the real-world testing of the PEV in close contact with people, one example that might contribute to the behavioral configuration of autonomous vehicles is an engagement process led by our colleague Edmond Awad. His Moral Machine study, a contemporary, gamified version of the classic Trolley Problem, enables internet users around the world to impersonate an autonomous vehicle in various Catch-22 scenarios where they are offered choices that all lead to casualties, with the typical scenario of choosing between saving oneself or saving others' lives through self-sacrifice. With over four million players/contributors from around the globe, the engagement enables researchers to dive into the moral norms, values, and preferences of people from different parts of the world and of different genders, ages, and even political preferences. In contrast to the more traditional approach of training AI systems that use limited—or even biased—data sets, this is an example of what Iyad Rahwan (2017) calls the "society-in-the-loop" approach to creating AI systems.

An early implementation of an autonomous robot that tries to adhere to social norms can be seen in another MIT research project that incorporated sensing and prediction of pedestrian conduct and crowd behavior in the robot's path planning, enabling what we might refer to as "socially aware navigation." The robot, in other words, dynamically adjusts its speed and path according to the surrounding conditions, which may include a single person or a group of people gathering or walking, with the goal of "traveling naturally among people and not binge intrusive... following the same rules as everyone else" (Chu, 2017). This type of incorporation of social awareness into the robot's path planning points to two important reminders for those who seek to develop AI applications for the transportation world. First, that the real world and street environment are inherently complex and the behaviors of users (pedestrians, cyclists, and motorists) can be unpredictable and are influenced by different local social norms. Second, in order for the robot to gain people's acceptance and trust for peaceful coexistence, or productive cooperation, the robots must first learn social conduct and behave according to the norm.

GETTING INFRASTRUCTURE READY

Shared bikes like the one we are testing through the PEV project will soon become on-demand, autonomous, and capable of providing logistics and urban services. The impact that they make will be greater if improvements are made to the road infrastructure. Across the world, city governments are quickly responding to the car-lite—or car-free—movement as a gesture to fight pollution and congestion. As cities build bike lanes, we are interested in the question of relevancy: is the city building bike lanes where they are needed? Are bike lanes improving the safety of cyclists and other road users? How might the new infrastructure support vehicles that are autonomous, or connected, to enhance the operations of new transportation solutions?

Using publicly available data, open source web maps, and a standard web browser, the City Science group at the MIT Media Lab has created a platform for visualizing supply and demand for cycling lanes in cities. The platform consists mainly of two layers of data visualized on the city map. The base layer, for showing supply, delineates the presence of bike lanes according to the data obtained from the city government, typically in the kml (Keyhole Markup Language) format, which can be easily converted to GeoJSON for visualization using web-based plug-ins. An additional layer shows demand by delineating the shortest path routes derived from the origin-destina-

FIGURE 2
NEIGHBORHOOD CHARACTERISTICS ASSOCIATED WITH INCOME MOBILITY

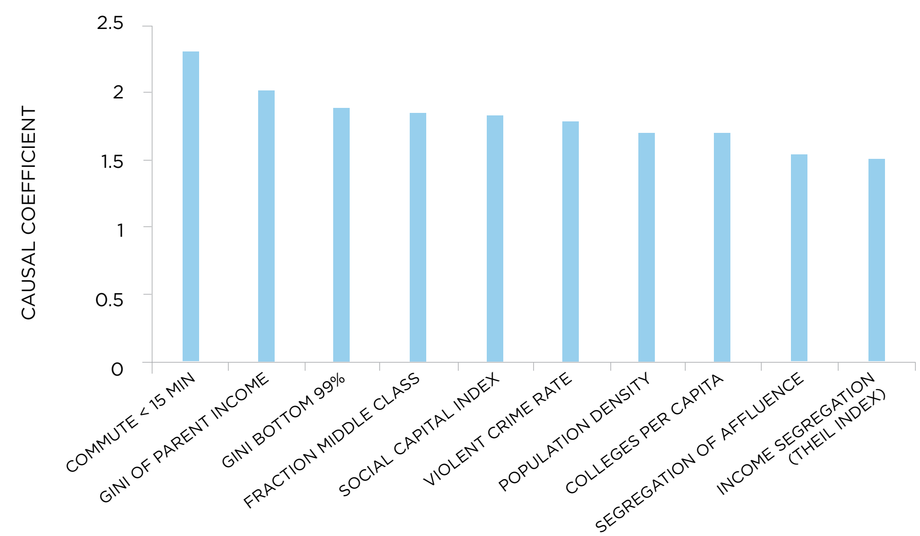

Source: Chetty and Hendren (2015).

FIGURE 3
AUTONOMY AND RIDE-SHARING COMPANIES

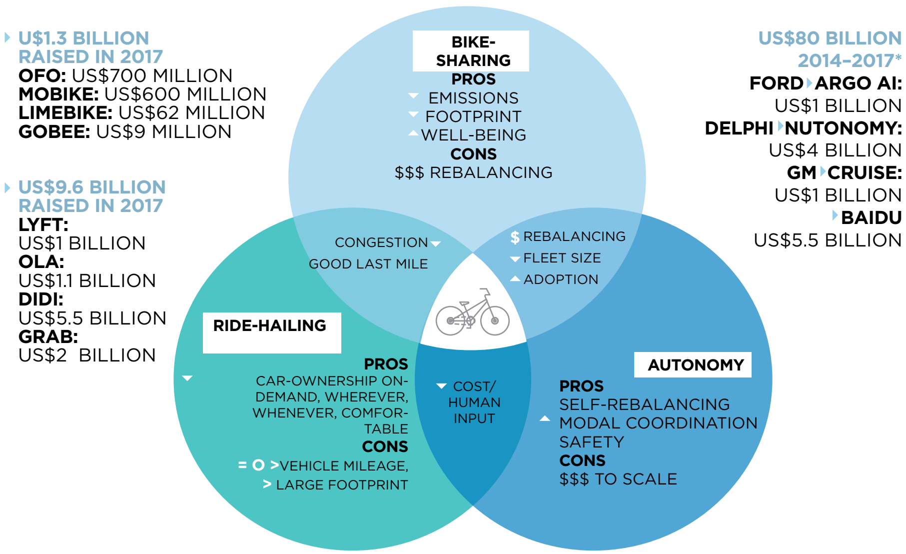

Source: Compiled by the authors based on data from the Brookings Institute

tion data shared by local bike-sharing operators, using the Open Source Routing Machine (OSRM)[1] developed by the open source community for generating navigation paths on OpenStreetMap[2]. Platform users thus have a map that compares road usage by shared bikes against the actual presence of bike lanes, and can quickly discern gaps that might exist between the two. Using data from Boston in the prototype version, one can observe that certain segments of the roads in the Seaport District are used by shared bikes more intensively than other parts during a period of time selected by the user. However, there are no bike lanes on these roads. In the case of Taipei, where the city government has recently published that it has built over 500 kilometers of bike lanes, we use the platform to separate and visualize the data into officially designated bike routes (road segments with mixed traffic but typically marked by a bike sign) and dedicated bike lanes. This enables analysts to compare the presence and absence of dedicated bike lanes with the locations where bike accidents have occurred, also using open data released by the city government.

The contrasts drawn in Boston between road usage and bike lane locations, or in Taipei between bike accidents and protected bike lanes enable various top-down and bottom-up actions.

As examples of top-down initiatives, the local public works or transportation agencies can use the platform to compare the gaps between different districts, and, in turn, prioritize funding toward where the gap appears to be greatest and bike lanes are needed most. Public agencies can also use such a visual supply-demand analysis as a starting point for further investigation and negotiation around issues such as road closures to automobiles, weekend street fairs, festivals, or actual road-use reassignment.

Bottom-up actions include local civil societies or elected officials quickly taking an evidence-based approach to lobbying for infrastructure budgets. As this platform is built for standard modern web browsers (e.g., Firefox, Chrome) using standardized data formats, open source maps, visualization, and navigation routing modules, it is easily scalable across different cities. Public agencies and community members can deploy their open data to inform infrastructure policies that support active mobility.

In addition to using open data to enhance cycling infrastructure for both existing and future bikes, we also anticipate that the design of streets and signage will play a critical role in the roll-out and "democratization" of autonomous vehicles. Recent advancements in autonomous technology have led car makers and tech companies to focus on reinventing the automobile by increasing computational capability and enhancing sensor systems. But due to strict road-safety regulations, this vehicle-centric, inside-out approach may take years to materialize, and when it does, restricting autonomy to selected vehicles will limit its impact on street safety and accessibility. To address this potential gap between technology and the public interest, through our Urban Tattoo project we are also investigating ways to offload often heavy computational requirements from the vehicle by creating affordable traffic signs and urban markers that can be read by both humans and machines (using computer vision). With the support of a new genre of smart urban infrastructure, we believe this "autonomy-lite" approach will soon allow lightweight autonomous vehicles, such as the PEV, to be widely deployed and navigate smoothly in most urban environments.

In this article, we have suggested three approaches that the City Science Group at MIT Media Lab is actively testing in collaboration with our partners around the world to help bridge the gap between the interests of industry and the public in the development of future transportation solutions. If the cars of the 20th century turn out to be the horses of the 19th century, we hope to see private-sector and institutional resources deployed toward methods that enhance the well-being and prosperity of people in the increasingly dense urban environment of the 21st century, such as reimagining bike-sharing using AI, cocreation, and open data.

NOTES
[1] http://project-osrm.org/.
[2] https://www.openstreetmap.org/.

REFERENCES

Bouchard, M. 2015. "Transportation Emerges as Crucial to Escaping Poverty." The New York Times. May 7.
Chase, R. 2018. "As Cities Grapple With Mobility Revolution, 10 Principles Emerge to Guide Them." The City Fix. February 14.
Chetty, R., and Hendren, N. 2015. "The Impacts of Neighborhoods on Intergenerational Mobility: Childhood Exposure Effects and County-Level Estimates." Harvard University and NBER.
Chu, E. 2017. "New Robot Rolls with the Rules of Pedestrian Conduct: Approach May Enable Robots to Move around Hospitals, Malls, and Other Areas with Heavy Foot Traffic." MIT News. August 29.
Clewlow, R.R. 2017. "New Research on How Ride-Hailing Impact Travel Behavior." Planetizen. October 11.
Duduta, N. 2010. "Are VMT and GDP Really Correlated?" The City Fix. June 8.
E360. 2017. "Transportation Replaces Power in US as Top Source of CO_2 Emissions." Yale Environment 360. December 4.
Ecola, L., and Wachs, M. 2012. "Exploring the Relationship between Travel Demand and Economic Growth." Report for the Federal Highway Administration. Washington, DC.
Fillin-Yeh, K. 2017. "Cities Can't Afford to Let 'Rogue' Bike Shares Run Wild." CityLab. June 18.
Foot, P. 1967. "The Problem of Abortion and the Doctrine of the Double Effect." *Oxford Review* 5: 5–15.
ITDP. 2013. *The Bike Share Planning Guide*. New York: Institute for Transportation and Development Policy.
Kerry, C.F., and Karsten, J. 2017. "Gauging Investment in Self-Driving Cars." Brookings. October 16.
Litman, T. 2018. *Evaluating Active Transport Benefits and Costs: Guide to Valuing Walking and Cycling Improvements and Encouragement Programs*. Victoria Transport Policy Institute.
Rahwan, I. 2017."Society-in-the-Loop: Programming the Algorithmic Social Contract." MIT Media Lab. July 24.
UN. 2014. "World's Population Increasingly Urban with More Than Half Living in Urban Areas." United Nations. New York, July 10.

CASE STUDY

Driverless Vehicles

Safety: The Primary Objective

37% OF TRIPS IN 2030 WILL BE IN DRIVERLESS VEHICLES

The automotive industry is one of the sectors that invest most in using new technologies in products and processes. The shift toward driverless vehicles, the inclusion of services within vehicles, the need to reduce traffic accidents, and ridesharing are some of the main factors driving this trend. PwC has estimated that 37% of the trips made in 2030 will be in driverless vehicles. Japanese firm Toyota ranks 10th on the global list of firms that invest most in R&D, which it spends US$8.8 billion on each year. In partnership with the Massachusetts Institute of Technology (MIT) and the University of Michigan, the company has established the Toyota Institute Research (TRI), which conducts research in the fields of artificial intelligence (AI), robotics, automated driving, new materials, and user experience to improve safety standards and develop intelligent vehicles. The TRI's communications manager, Rick Bourgoise, explains the main projects that the company is working on and how AI is impacting the industry and the future of mobility as a service (MaaS).

What challenges is the automotive industry facing?

There are three main components to automated driving: perception, prediction, and planning. Perception is what the automated vehicle sees through its cameras and sensors, while prediction is the ability to anticipate what other vehicles will do. Planning is the process of making decisions to define what the automated vehicle will do in response to a given situation. We think that we have a firm handle on perception and planning, although there are still some significant shortfalls, such as achieving good perception in the middle of a blinding snowstorm. We understand these technical challenges and we are working hard to move beyond them. AI-based prediction is still the big question. Human drivers are very good at predicting behavior on the road, so machines need to become even better at predicting and anticipating this. We have a lot of work ahead if we want to develop AI that can handle the infinite number of scenarios needed to make fully autonomous driving a reality. Automated driving is relatively easy to achieve in empty spaces, but the real challenge is handling real-world situations.

What are the main innovation projects that Toyota is working on?

The TRI is researching driverless technology for Toyota. We recently unveiled our next-generation driverless research vehicle, Platform 3.0. Our work has really come to maturity through this platform, which includes a rich sensor package that makes it one of the most perceptive driverless test cars out there. We are currently exploring two other approaches to vehicle automation, which are called Guardian and Chauffeur. Chauffeur is Toyota's trial version at the SAE 4/5 autonomy level, which means that the automated system carries out all driving tasks. Guardian, on the other hand, uses the automated driving system as a safety net for human drivers to prevent accidents.

What is the outlook for the sector over the next 10 years?

Over the next decade, we are likely to see an increase in the number of vehicles with high levels of driver assistance technology, mainly in urban areas with established operational parameters (such as geofenced areas or ideal weather conditions). We think that the rollout of driverless vehicles will initially be for MaaS applications. Service vehicles can generate income almost around the clock and can offset high-technology costs faster than privately owned vehicles. We don't anticipate the disappearance of vehicles driven by humans in the foreseeable future, and maybe not even over the next few generations. Advances in the three Ps that I mentioned above (perception, prediction, and planning) are fundamental to the seamless coexistence of human and automated driving.

What are the benefits of an automated car?

Safety. Advanced vehicle technology has already proved that it can save lives, and we are striving to take this even further. The Guardian and Chauffeur applications that we are developing are bringing Toyota closer to its global vision of zero deaths in traffic accidents. Our aim is to build a car that is incapable of causing an accident. Sensor technology is another key piece in this puzzle and we are seeing dramatic improvements in LIDAR technology, which detects the relative position of objects in the vehicle's environment. We are also exploring technology that can monitor drivers to detect if they are drowsy or distracted.

Do you think that people will accept these vehicles?

We'll need to educate people. Creating trust is key. Recent studies in the US have shown that a large percentage of the general public is skeptical about these technologies and would be afraid of traveling in a driverless vehicle. There needs to be more emphasis on educating the public about the safety benefits that this technology could bring and showing people how these vehicles work.

Port Infrastructure in 2035

Luis M. Ascencio
Latin American and Caribbean Economic System (SELA)

Rosa G. González-Ramírez
University of the Andes

ARTIFICIAL INTELLIGENCE IS GAINING MOMENTUM AS A TOOL FOR ACCELERATING THE PRODUCTIVITY OF PORTS BOTH ON LAND AT SEA. SENSORS NOW ALLOW CARGO STATUSES TO BE MONITORED ALONG THE ENTIRE GLOBAL SUPPLY CHAIN AND AUTOMATED GUIDED VEHICLE (AGV) SYSTEMS MOVE CONTAINERS AUTONOMOUSLY. THIS ARTICLE EXAMINES THE FUTURE CONVERGENCE OF INFRASTRUCTURE (MACHINERY) AND INFOSTRUCTURE (DATA AND CONNECTIVITY), WHICH IS BEING MADE POSSIBLE BY THE MATURITY OF DIGITAL PLATFORMS.

The classic approach to management in the shipping and port industry centered on operations within port terminals themselves. Digitalization has shifted the focus to the entire logistics chain. The digital economy has brought new opportunities for the industry to increase its productivity, efficiency, and sustainability (Heilig, Schwarze, and Voß, 2017). As a consequence, data systems and technology have come to play an essential role in guaranteeing competitiveness, as they facilitate communication between all stakeholders through digital exchanges of data, automation, and streamlining processes (Heilig and Voß, 2016).

This article discusses a vision of how the shipping and port industry will evolve and the actions that are needed to facilitate the adoption of new management strategies and disruptive technologies. It also makes recommendations to help both private stakeholders and governments in Latin America achieve this. The article is structured as follows: the second section provides a summary of the state of the art of technologies and management capacities in the digital age, describing the main areas of development in the field of emerging ICTs. The third section contains an overview of how the shipping and port industry has evolved as technology, infrastructure, infostructure, and management models have converged. The fourth section puts forward an emerging management model that contemplates the inclusion of disruptive technologies and business ideas known as lean and green port logistics. The fifth section examines the relevance of public policies for the development of the shipping and port industry in the digital economy, presents some key components of this, and cites examples of concrete action to achieve this in Latin America. The final section puts forward some general conclusions and recommendations for furthering a regional agenda.

NEW TECHNOLOGICAL AND MANAGEMENT CAPACITIES

We are living through a time of transition within the digital age, one in which technology—the outcome of an evolving, disruptive process—is impacting different sectors of the economy and our social and personal lives (Tapscott, 1995). Experts argue that by 2020, we will see the consolidation of a digital economy that is based on connections between people and their communications with each other (using mobile phones) and businesses, generally through online means of exchanging documents and service orders.

The path the digital economy has taken up to now has been fraught with complexities. Once the market had

CLOUD COMPUTING IS VITAL TO SCALING UP COMPUTING RESOURCES

moved beyond the dot-com bubble of the late 1990s and the negative knock-on effects of this on credibility and market confidence, better internet-based solutions with a more pragmatic commercial focus arose. Narrowband internet made a significant impact on society through the mass adoption of e-mail and functional websites. Then, between 2005 and 2010, broadband began to facilitate higher-speed data transmission and established itself as the information highway by offering global platforms for people, companies, governments, and ultimately, machines and things to interact, communicate, work together, and access information.

The economy is currently said to be going through a transition stage because the phenomenon of online devices is on the rise, which poses enormous challenges around the use of more intelligent and better-connected products that provide new functionalities, are more reliable, and make more efficient use of resources.

In a recent study, the consultancy and analytics firm Gartner gave an answer to one of the key questions about this change: how many online devices will there be in the future? In response to this, the company published a report which predicted that there would be over 8.4 billion objects throughout the world by the end of 2017, increasing to 11 billion in 2018, and 20.5 billion in 2020. The consumer segment would be the largest user, followed by cross-industry business and vertical-specific business (Ferrer Caballero, 2017). Futurists say there will be 50 billion online devices by 2030.

There are five main areas of development in the field of ICTs that are impacting the speed at which people, companies, and devices are adapting to the digital economy (SELA, 2017). The first of these is cloud computing. This technology supplies on-demand IT resources for computing and storing data online, which facilitates consumption-based payment models. The second key area is big data analytics, which is also known as data science. This is connected to techniques for handling large databases such as the Hadoop framework, an open-source Apache project for the distributed processing and storage of large data sets across clusters of computers. The real-time availability of data and mechanisms for processing it, finding patterns in it, and identifying ways of using it to support decision-making are all extremely valuable for industries. It is particularly significant for the shipping and port industry because it is a highly complex sector where multiple stakeholders, documents, and information sources need to be coordinated, which generates a huge number of transactions.

The third key area is the Internet of Things (IoT), which is a network of physical devices that are connected to the internet through embedded sensors (MHI, 2017). This technology, in combination with others such as data analytics, helps make information available to decision-makers in real time. This is particularly relevant to shipping and ports: the extreme complexity of the industry means that the ability to make predictions around variables of interest is particularly valuable, as is being able to track and trace cargo. For example, being able to predict the arrival pattern for trucks at a sea terminal would allow the operations planner to decide which resources to assign to each task (handling ships or dispatching containers to trucks).

The fourth area of development is cryptocurrencies and blockchain, which have revolutionized the way business is done. Cryptocurrencies are digital means of exchange, the first of which was Bitcoin. Blockchain is a distributed database made up of chains of blocks of records which are designed to prevent their being modified whenever a new record with a reliable timestamp is published and linked to a previous block. This sidesteps the need for intermediaries in transactions and makes information available in real time. This is why the shipping giant Maersk has signed agreements with IBM to develop more efficient global trade mechanisms using blockchain (White, 2018).

The last area of development that we will focus on here is artificial intelligence (AI). According to López Takeyas[1], this is a branch of computer science that studies computational models to carry out activities usually performed by human beings. The concept first arose in the 1950s and includes a range of methodologies, such as machine learning, visualization, and advanced algorithms. These methodologies are also commonplace in the field of data science or data analytics. The main difference between them is that AI is oriented toward carrying out activities that are essentially human, such as reasoning and understanding behavior.

There have already been documented examples of the implementation of AI in the shipping sector. One example is Kuznetsova, Spellman, and Jumma's (2018) article "How Artificial Intelligence Can Power Growth and Opportunities in Global Container Shipping." The authors are part of the executive team at INTTRA, a digital B2B platform that provides information and solutions for the shipping industry. It was founded in 2001 as a joint venture between a group of shipping companies to create a neutral platform for booking cargo space and exchanging information between shippers and carriers. In Kuznetsova et al. (2018), the authors refer to the disruptive technologies discussed above

THE PORT OF ROTTERDAM'S VISION

The Port of Rotterdam's new vision is built on two main pillars: the Global Hub and Europe's Industrial Cluster. The port envisions close cooperation and "partnerships between businesses, government agencies, and knowledge and innovation institutions, which will result in a high-quality labor market, living environment, and excellent accessibility. [...] Adaptability is the key word." By combining the features of the Global Hub and Industrial Cluster, the port envisions that the following factors will define success:

-investment climate;
-use of space;
-accessibility;
-shipping;
-laws and regulations;
-the city and region;
-work;
-knowledge development and innovation;
-environment, safety, and living environment;
-Europe (regional division).

Source: Port of Rotterdam (2018).

and argue that in the coming years, these technologies will lead to massive improvements in the way that products and services are shipped. In other words, they connect physical cargo flows with information flows.

Although AI is understood as being tools that perform essentially human activities, Kuznetsova et al. (2018) point out that rather than replacing human workers, this technology is currently being used to support, facilitate, and help them go about their jobs. This may also be because AI systems continue to produce errors in their predictions and visualizations, so humans cannot yet be replaced altogether—indeed, they remain a critical factor when handling exceptions or outliers and carrying out highly complex tasks. Although many procedures in the shipping and port industry can be standardized and automated, there are a huge number of exceptions. Consequently, rather than replacing jobs, AI will allow workers to focus their energies on these exceptions and on making complex decisions, while the more repetitive tasks are automated. The use of machine learning and data analytics and the possibility of storing large volumes of data on demand mean that information is available for making decisions in real time.

A tool that INTTRA has implemented to predict demand for container transportation and assign space is one of the different AI applications being used in the shipping and port industry in combination with the other technologies described above. INTTRA's tool also saves information from clients' past bookings, so all the relevant information is automatically filled out for them when they make a new one. Another area in which INTTRA foresees potential for AI is the Harmonized System, the international standard used to classify merchandise being traded. Errors entering these codes can lead to delays that AI applications could prevent. Intelligent containers are another development that the INTTRA executives discuss in their publication (Kuznetsova et al., 2018). This is a system that uses AI to track and monitor containers remotely. For example, it can be used in reefers (refrigerated containers) to monitor temperature. The aim is for the system to only call on human personnel in exceptional circumstances.

THE EVOLUTION OF THE SHIPPING AND PORT INDUSTRY

Ports and the shipping industry are fundamental parts of the global transportation system. The acceleration of global trade flows is opening up a range of possibilities and challenges for countries, government agencies, the private sector, and the port sector in particular, whose infrastructure, infostructure, and management models are truly beginning to converge.

The increased flow of goods and the growing importance of this in supply chain competitiveness have prompted many ports in developed countries to update their commercial strategies with a view to the future. The Port of Rotterdam is one such example: in 2011 it published its "Port Vision 2030" (see box 1).

Based on the current situation and what is envisioned for the shipping and port industry in the future, figure 1 presents an overview of how the industry will evolve, which is analyzed in view of the convergence between info- and infrastructure (y-axis), approaches to management (x-axis), and time (secondary axis). Physical infrastructure is the core component of the y-axis, to which infostructure was then added, based on the exchange of digital data, IoT, and AI.

The main feature of stage 1 was reducing costs and increasing productivity for both the shipping industry and ports. This 20-year period (1990–2010), saw bankruptcies and mergers and acquisitions in the shipping industry; an increase in the capacity of container ships; the consolidation of global terminal operators (GTOs) due to the drive to reform state port models; major investments in new ports; and the adoption of better, new, more specialized high-productivity port equipment. In sum, this was a stage in which physical infrastructure was a key factor for both businesses, which continued to function separately, with limited integration.

Stage 2, which we are currently in the midst of (2010–2025), focuses on integration, which has been facilitated by the digitalization of the shipping and port industry on a larger scale and a significant change in the management model, whereby the concept of supply chain management, which was historically linked to the textile, automotive, electronics, and retail industries, is now viewed as a tool that will accelerate both productivity in ports' land- and sea-based operations. This phase is characterized by business partnerships between shipping companies to share space and bring down costs, and the global emergence of port community systems (PCSs), which has led to the consolidation of tech firms providing such services. This phase has also seen the signing of trade and transportation facilitation agreements between countries through the hard work of the United Nations Centre for Trade Facilitation and Electronic Business. Another feature of this stage is the widespread use of the internet in the vast majority of business processes. It has also been marked by a new enabler called orgware, which explores new forms of public–private governance in connection with the port industry.

FIGURE 1
STAGES IN THE EVOLUTION OF THE SHIPPING AND PORT INDUSTRY

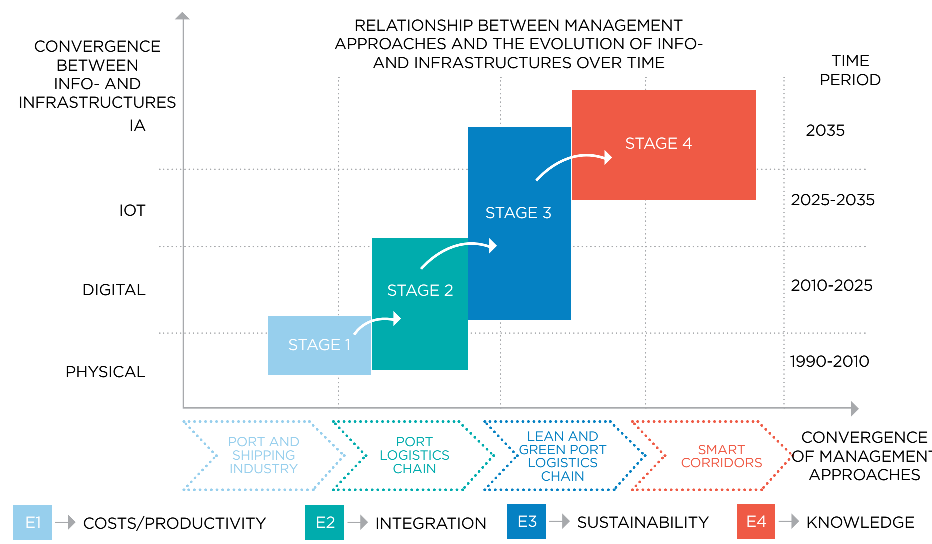

Source: Compiled by the author.

Stage 3 (sustainability) will come to fruition between 2025 and 2035, and its main feature will be the convergence of infrastructure (mainly machinery) and infostructure (data and connectivity) as IoT platforms come to maturity. These platforms are currently being concept-tested for use in different areas of the shipping and port industry. For example, pilot tests for containers and sensors are taking place to monitor the status of freight along the global supply chain; AGV equipment is being used to move containers around ports autonomously; and testing is being carried out to integrate data on land-based hinterland freight transportation with ports (synchromodality). This management approach is known as lean and green port logistics as it strives to provide environmentally responsible services that make much more effective use of assets. IoT platforms will operate using an online blockchain-based communications protocol, which will enable safer, lower-cost operations without the need for intermediaries. There will also be an exponential increase in data volumes, and big data companies will come into their own as major new stakeholders in the global transportation services ecosystem. Stage 3 will clearly be a time of transition.

Stage 4 (knowledge) visualizes a scenario in which pre-established communication protocols between the machines that governed IoT will have given way to large-scale implementation of AI. This will allow complex machines and cybernetic systems to control a series of service areas that are currently being coordinated by human beings or hybrid models (humans assisted by machines). Autonomous freight vehicles could be implemented outside port areas and may include ships, trucks, containers, drones, and other parts of the transportation system. These vehicles will use their own languages to make decisions and ensure forms of mobility that are free of errors, inefficiencies, waste, unnecessary costs, and, in many cases, people. The port management approach will shift toward one that is based on smart corridors, real networks of autonomous transportation that are efficient, resilient, and take a hybrid approach to the movement of goods and persons (see box 2).

SMART PORTS

Smart ports promote the adoption of disruptive technologies that enable better planning and management of the port logistics chain. This concept stretches beyond the boundaries of port terminals themselves and into the surrounding communities. The Port of Hamburg was one of the first to use the term, through its smartPORT program, which was implemented by the Hamburg Port Authority. The use of sensor technology, in combination with data analysis and forecasting and information systems, has led to substantial improvements in the efficiency of its operations, which benefits both the port itself and the surrounding area. Through its smartPORT philosophy, the Port of Hamburg has achieved sustainable economic development and has maximized benefits for stakeholders in the port community and the surrounding area.

Source: Hamburg Port Authority (2018).

AN EMERGING MODEL

As we discussed above, the process of integrating the shipping and port industries is constantly evolving, driven by ever-greater flows of goods. These industries' aims and strategic objectives can only be reached through the permanent adaptation of management models using available physical and digital technologies. The main feature of the new management model known as lean and green port logistics is the introduction of disruptive technologies that have now moved beyond the prototyping and lab testing phases and are being included in businesses' new service requirements.

Many the world's leading port systems are already using this emerging management model which, as figure 2 shows, combines the factors needed to reach the strategic goal of sustainability, balancing out the economic, social, and environmental factors affecting the shipping and port industry. Each component in this emerging management model is explained by a series of factors that have been maturing both inside and outside the industry. Outside the industry, disruptive technologies such as IoT, blockchain, big data, and cloud computing are being facilitated by the growing reliability and availability of mobile internet services. This is improving data processing capacities and giving rise to new technological infrastructure services. Furthermore, processed data is making a huge impact on the efficiency of the extended port logistics model, which is known in Europe as the synchromodal model.

Understanding how large databases are organized is important. With data flowing through technologies such as blockchain, it is possible to make payment information, documents, and shipping orders more secure within the port logistics chain while improving the interoperability of information systems between business stakeholders, communities, and domestic and international foreign trade and logistics platforms. Likewise, the real-time exchange of data between stakeholders is helping to refine decision-making around the use of critical assets such as storage capacity in ship holds, cranes and warehouses at ports, and trucks and railways. This is facilitating the synchronization and efficient integration of intermodal transportation systems.

> **BLOCKCHAIN**
> TECHNOLOGY WILL ENABLE MORE SECURE OPERATIONS AT A LOWER COST

Internally, ports being managed in this way have intensified their governance models, which have been transformed from a mere exercise in management by the state or port authority to an effective form of comanagement with the private sector and business and innovation networks in the port city. At this stage of development, formal legal port logistics communities are the cornerstone of governance, which has allowed them to implement coordinated actions to increase the competitiveness of ports in terms of the quality and safety guaranteed to users, implement new environmental standards for clean production, integrate into surrounding areas and port cities, and, finally, become a driver for trade facilitation through members opting to become part of authorized economic operator (AEO) programs.

US$ 50 BILLION
THE INVESTMENT NEEDED IN LATIN AMERICA'S PORTS

Finally, an essential part of adopting new technologies and disruptive business models in the shipping and port industry is human capital. Education and training focus increasingly on soft skills such as being customer-focused, understanding value chains, being able to adapt to constant change, and working in an environment of collaborative competition. These new skills are being taught jointly by port communities and schools in developed countries such as Singapore, Germany, the Netherlands, Belgium, and France, which are blazing the trail for other countries that still have gaps to bridge if they are to fully apply this emerging management model.

PUBLIC POLICIES

As SELA (2017) argues, the use of technologies such as cloud computing, data analysis, IoT, AI, and blockchain requires governments to make major efforts to outline public policies that strengthen their digital ecosystems and encourage various sectors to adopt new technologies, particularly the port and shipping and foreign trade sectors. The Digital Agenda for Latin America and the Caribbean (eLAC2018) was designed to strengthen the region's digital economy. It puts forward a strategy for 2018 based on governments committing to designing programs that promote access to and use of digital technologies. Work began on this ECLAC-led initiative in the year 2000 and various work areas have been developed to date, including the Ministerial Conference on the Information Society in Latin America and the Caribbean. The sixth conference is planned for 2020 (eLAC2020).

Those ports that have made headway on implementing emerging management models have been able to combine and take advantage of the digital economy in their social and institutional environment (external context) and through the intelligent management of the competitiveness variables that are inherent to the port and shipping business (internal context). Latin American ports are still lagging behind in terms of technology use and the implementation of new management models. As part of an international technical cooperation initiative, in 2014 the Latin American and Caribbean Economic System (SELA) and the Development Bank of Latin America (CAF) launched the Latin American and Caribbean Network of Digital and Collaborative Ports, the main objective of which is to contribute to disseminating and transferring knowledge and information on new management trends, the digitalization of public and public-private operations, and logistical competitiveness within ports. The program currently includes 26 port systems in 13 countries from every subregion in the Americas.

One of the mainstays of the program is the promotion of port logistics communities from the region and connecting the different public and private stakeholders involved in them, along with sectoral regulatory organizations and universities or research centers.

While governments debate how to apply new and better public policies to the transportation sector, the market continues along the path to meeting the growing demand for freight transportation by developing suitable infrastructure to meet these needs. A recent study by CAF (Arroyo Crejo, 2018), "Analysis of Port Investments in Latin America and the Caribbean with a View to 2040," shows that by the date in question, the region would need to have invested nearly US$50 billion to meet the estimated demand of 150 million 20-foot equivalent units or containers (TEUs). During this period, Latin America will go from having six ports that can currently move 2 million TEUs to 20 ports that can do so.

PROMOTING THE REGIONAL AGENDA

The new technologies reaching our homes, businesses, and governments are changing the ways that we work and interact, presenting new opportunities for us to make progress as a region and develop stronger, more resilient communities of interest. Emerging technologies like IoT and AI, among others, may help people and organizations to become more efficient and achieve more effective results that translate into greater productivity and economic growth.

Contrary to expectations around automation and the use of emerging technologies such as AI, it is important to stress that the aim of these tools is not to reduce the workforce. There are still huge numbers of exceptions that these tools are unable to handle alone and which require human attention. In highly complex industries like shipping and ports, the aim of these technologies is provide support for workers as they go about their activities and make decisions, so that simple tasks can be automated while executives focus on dealing with more complex tasks or ex-

FIGURE 2
TOWARD LEAN AND GREEN PORT LOGISTICS
TECHNOLOGIES, INDUSTRIES, AND DISRUPTIVE OPERATIONS

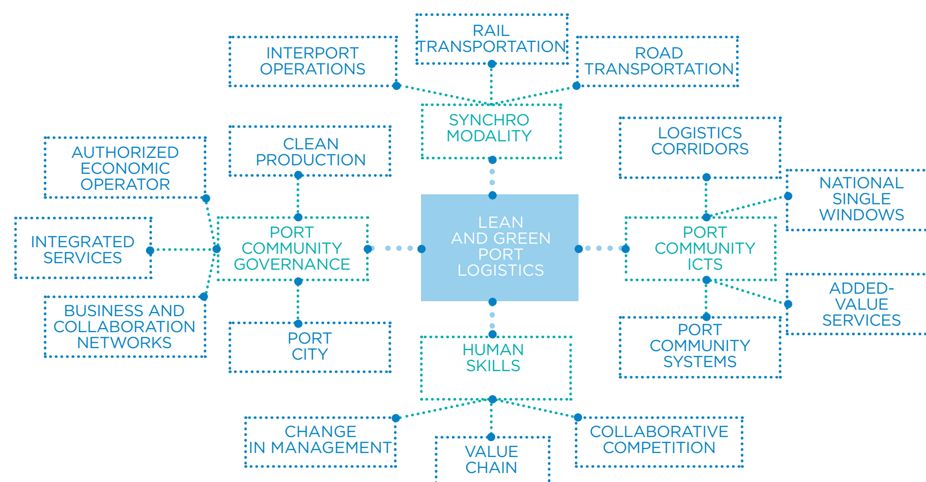

Source: SELA–CAF Network of Digital and Collaborative Ports.

ceptional circumstances.

To ensure that firms adopt these technologies and that people can benefit from digitization, it is essential to provide top-down directives while fostering bottom-up collaboration.

In the 20 years between now and 2040, Latin America's shipping and port sector will see a threefold increase in overall freight traffic volumes (import and export containers). This will imply an estimated investment of US$50 billion to ready 20 ports in the region for handling over 2 million TEUs each per year.

To take on this challenge, the sector needs to ensure that logistic integration and sustainability are part of its management strategies. The use of emerging technologies will be essential to delivering maximum value in response to growing needs for better, more connected flows of cargo and foreign trade.

Evolving technology and the need for change within management models in the shipping and port industry mean that governments and international technical cooperation have an essential role to play in articulating agreements around the digital economy and the ways that this is impacting specific industries. Implementing a strategic regional agenda to facilitate the development of smart ports will help countries build their capacities for integrating into logistics chains and economic integration corridors; promote shipping traffic; develop efficient investments in infrastructure; increase their competitiveness and operational efficiency; and, finally, become institutionally, environmentally, and economically sustainable.

Regional forums such as the Network of Digital and Collaborative Ports are the ideal place for governments, the industry, and academia to establish appropriate strategies to address the challenges posed by the digital economy.

NOTES

[1] López Takeyas, B. "Introducción a la inteligencia artificial." Available at: http://www.itnuevolaredo.edu.mx/takeyas/Articulos/Inteligencia%20Artificial/ARTICULO%20Introduccion%20a%20la%20Inteligencia%20Artificial.pdf

REFERENCES

Arroyo Crejo, F. 2017. "Desarrollo portuario en América Latina y Caribe. Horizonte 2040. [Reflexiones para debatir]." Paper presented at the II Meeting of Port Logistics Communities in Latin America and the Caribbean. CAF-CELA. Cartagena de Indias, Colombia.

—. 2018. Análisis de inversiones portuarias en América Latina y el Caribe al horizonte 2040. Caracas: CAF.

Ferrer Caballero, C. 2017. "8.400 millones de dispositivos estarán conectados a Internet a finales de 2017." World Economic Forum, February 20.

Heilig, L., and Voß, S. 2016. "Information Systems in Seaports: A Categorization and Overview." Information Technology and Management 18 (3): 179–201.

Heilig L., Schwarze, S., and Voß, S. 2017. "An Analysis of Digital Transformation in the History and Future of Modern Ports." In: Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS). Big Island, Hawaii: HICSS.

Hamburg Port Authority. 2018. "smartPORT. The Intelligent Port." Available at: https://www.hamburg-port-authority.de/en/hpa-360/smartport/.

Kuznetsova, I., Spellman, P., and Jumma, K. 2018. "How Artificial Intelligence Can Power Growth and Opportunities in Global Container Shipping." INTTRA.

MHI. 2017. Material Handling & Logistics U.S. Roadmap 2.0. MHI.

Muruzábal, C. 2018. "Para que América Latina prospere en la era digital, primero debe enseñar a las mentes, luego las máquinas." World Economic Forum, March 15.

Port of Rotterdam. 2018. "Port Vision 2030." https://www.portofrotterdam.com/en/the-port/port-vision-2030.

SELA. 2017. "Comercio y competencia en la era de la economía digitalizada." Working Paper No. 2-17. San Salvador: 7th Annual Meeting of the Working Group on Trade and Competition in Latin America and the Caribbean (GTCC), SELA.

Tapscott, D. 1995. The Digital Economy: Rethinking Promise and Peril in the Age of Networked Intelligence. New York: McGraw-Hill Education.

White, M. 2018. "Digitizing Global Trade with Maersk and IBM." IBM. January 16.

TECHNOFACILITATION

Exponential technologies and the current digital revolution offer a range of possible ways for helping the organizations that oversee foreign trade in Latin America and companies that are active in this sector. Artificial intelligence (AI) using big data analysis has various potential uses in trade facilitation. For example, Brazil has developed an AI system to detect different types of import/export fraud. AI is also being used to draft more efficient contracts, facilitate access to commercial credit, and save time complying with requests for pre- and postsales support. A system operated by 3CE automates the classification of products using the Harmonised System, reducing the time spent on this process and minimizing costly classification errors. However, regardless of what new technologies are capable of, the factors that are most hampering progress on the region's trade agenda are the same as always: the role of institutions, human resource capacities, and access to financing. Furthermore, in a context of large digital divides—Internet use rates in the region range between 30% in El Salvador and Honduras and 70% in Argentina and Chile—the use of technology in the world of trade could end up promoting highly uneven progress on trade facilitation, which would be incompatible with the idea of greater regional integration.

70%
the progress LAC has made on the WTO trade facilitation agenda

2.5 HOURS
the time a foreign trade transaction takes using blockchain

40%
the reduction in container selection errors using augmented reality

Major Challenges in Implementing Trade Facilitation Measures

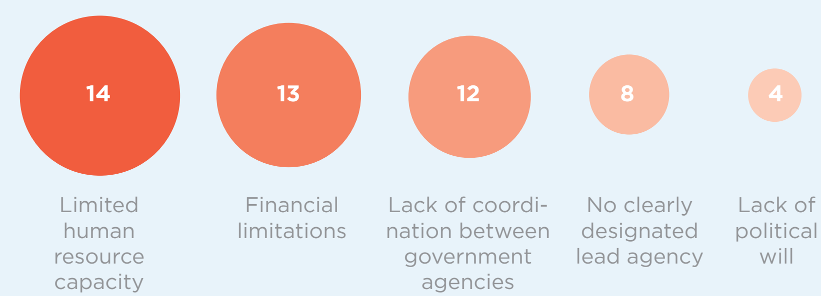

Source: INTAL/IDB based on a United Nations survey of LAC countries.

CASE STUDY

Robotic Logistics

Efficient Inventory Management

Before we start sharing the road with driverless cars, we're likely to run into robots roaming shops or warehouses. The artificial vision technology that enables machines to recognize spaces and sense objects can be used for more than just navigation: it also helps them identify objects, automate tasks, and manage inventories.

Fellow Robots is a Mexican company that specializes in using autonomous robots to perform inventory and logistics work. Marco Mascorro, CEO of the Silicon Valley–based firm, has won prizes from *Forbes and the MIT Technology* Review for his innovations. In this interview, he explains how robots cruise the aisles of stores or warehouses, generate a map of the facilities, and recognize products displayed on the shelves. This keeps the digital inventory up-to-date and brings down costs, which can represent as much as 30% of the product price.

How might AI affect the logistics sector?

Robotics and AI are growing exponentially, but it's still very early days. At present, both are being implemented for basic tasks, but as technology and computing power develop and cloud computing costs come down, they will have a huge impact on the logistics and warehousing industry, where processes will be optimized extremely efficiently. At the same time, the amount of data that is generated will increase, and human decisions will become more sophisticated.

How much more efficient can warehouse or store management become?

These solutions are being implemented at a rapid pace. The advantage of many of them is that they live in the cloud, which enables them to become instantly global. This impacts distribution systems, transportation systems, and time to final destination. In the case of robotics, it helps us capture data on warehouses, shops, and other systems efficiently and continuously.

What can we expect in this field in the coming years?

The convergence of technologies like AI, robotics, cloud computing, and shrinking costs have enabled these technologies to be developed and implemented and to grow and spread. These systems are already in use in various areas of logistics, and in the short term we will see more of them being implemented in more traditional parts of the industry. However, in the coming years, more AI will be used for purchase order systems, inventory prediction, and to optimize distribution and transportation systems.

FAR FROM THE OECD AVERAGE

At present, logistics costs in Latin America range between 18% and 35% of final product prices, far higher than the 8% average for OECD countries. For SMEs, this percentage sometimes climbs above 40%. The region needs to double its investment in infrastructure to put itself on a par with that of developed countries, where more competitive infrastructure reduces companies' logistics costs. According to Mascorro, technology could play a part in closing this gap.

What does a company need to be able to implement and take advantage of these solutions?

These technologies are more accessible than ever and they cost much less than they did just a few years ago. The advantage today is that these systems can be implemented in existing environments without the need to for changes to infrastructure. This is another of the benefits: this technology can be flexible, learn to operate in a new environment under radically different parameters, and improve processes there. At present, we have implemented these systems in both large and small firms with minimum investment and the impact has been positive from the first day of operations.

What countries are you currently operating in?

Mainly in the United States but also in Japan. However, we think that the size of the Latin American market represents a huge opportunity.

Will automation lead to a loss of jobs in the logistics sector?

The aim of our technology has always been to complement human labor as efficiently as possible. Technology provides us with the data tools we need to make better decisions. Just as mobile phones have done in our everyday lives, these systems provide a way of accessing information, which never been as abundant or helpful as it is now.

Toward an Intelligent Form of Mining

John Atkinson
Adolfo Ibáñez University

ARTIFICIAL INTELLIGENCE IMPACTS PRODUCTIVITY AND BUSINESS MODELS IN DIFFERENT INDUSTRIES. IN MINING IT PLAYS A PART IN OPTIMIZING DIFFERENT STAGES IN THE MINING CYCLE, INCREASING EFFICIENCY, CUTTING COSTS, AND INCREASING WORKER SAFETY.

The mining industry is going through a period of intense change, so the ability to innovate and improve operations has become essential. Recent studies have estimated that true innovation will drive the next wave of gains in productivity and financial growth in the sector (Duddu, 2014; FN Media Group, 2017).

Fortunately, emerging automation technologies that will enable innovations in the mining industry are ready to change the way that mining companies operate and help them evolve toward new business models. One model that is set to grow in the short term is intelligent mining, which focuses on delivering solutions using information technology, process automation, and robotics. The emphasis will be on productive processes that are inherently risky for human operators due to the nature of mining itself (Durrant-Whyte et al., 2015). Mining has gone from being a process that revolved around physical work to an industry that makes intensive use of large-scale machinery, which opens the door to specific applications of knowledge and technology that will improve productivity and safety at work while also making less of an impact on the surrounding environment.

The use of automation in mining and decision-making processes is nothing new (Karakaya, 2017). One traditional motivation for this has been the systematization of routine tasks that are costly when carried out by human operators, even when they perform them well, which is often not the case. However, many complex production problems are due to the fact that some tasks cannot be performed efficiently even when using experienced human workers. For example, many tasks in the mining cycle produce huge quantities of data from different sources (GPS, sensors, images, etc.). Yet humans are unable to process all this data, analyze it, and make real-time decisions based on it (Habrat and Lisowski, 2015). Automated haul and transportation systems are another application: some are operated remotely or have a certain degree of autonomy using geolocating systems of some sort, but the systems currently in operation do not operate entirely on their own. Activities are usually planned semi-automatically, due to their complex, dynamic nature, and this prevents predictions from being made and events from being anticipated.

The solution to these problems is clearly not about mere traditional automation but rather about designing intelligent systems that can analyze situations and make independent decisions (without human intervention), efficiently exploring millions of alternatives and learning from experience. It is hoped that this type of technology will not only automate specific tasks but will also perform better at it than humans.

As in many other areas of production, advances in AI are developing quickly in the mining sector, changing the way it is structured and transforming the way that complex tasks are carried out, be they everyday activities or industrial ones, which could have unprecedented impacts on the global economy (Purdy and Daugherty, 2016; Rao and Verweij, 2017). Generally speaking, mining executives have to make many critical operational decisions while also complying with day-to-day obligations such as complying with safety standards and meeting production targets. These decisions usually involve offsetting complex factors (e.g., operating costs vs. production costs in processing plants), and relying on forms or simple criteria can lead to suboptimal solutions, so taking advantage of AI can help to avoid this trap (Karakaya, 2017).

The large quantities of available data and potential data sources and the increase in low-cost sensors and intelligent devices means that AI can be used to solve problems that were impossible to solve as recently as a few years ago. Furthermore, new intelligent data analysis algorithms and increases in computational power (e.g., cloud computing and high-performance computing) have made AI applications much more accessible in many productive spheres.

This article examines how AI can improve the productivity and efficiency of the mining sector. To this end, it describes a range of mining operations where AI has been successfully applied around the world and the projected impact of these changes.

THE AUTOMATED MINING CYCLE

A study from PwC (2017) shows that the 40 largest global mining corporations have market capitalizations of US$748 billion. The industry as a whole experienced a recession in 2015. It has since recovered, due to increases in commodity prices, and the sector has made a significant effort to improve efficiency at all levels. Small changes in automation that improve speed, performance, and efficiency usually set a profitable operation apart from an unprofitable one.

The main advantage of automation in mining is making the safe and efficient monitoring of any industrial process a viability (Durrant-Whyte et al., 2015). The nature of many mining tasks makes them hard for human experts to monitor because they entail many variables that are difficult to oversee and keep track of due to the speed at which they change, the complex relationships between them, the workplace risk for an operator, the degree of precision required, and so on. Innovation is fundamental for including factors such as mobility, data analysis, and intelligence in different mining operations. The application of automation technologies in mining facilitates the integration of digital intelligence into mining processes, which in turn enables new forms of interaction and ways of using information for real-time decision-making (FN Media Group, 2017).

AI is a major technological trend that is revolutionizing different production processes and the economy as a whole and it may have a significant impact on digital mining. This field of computer science creates intelligent computing systems that work and react like humans to solve complex problems (Russell and Norving, 2015). This spans a broad spectrum of activity, from voice recognition and visual perception in robotics to machine translation and decision-making of the sort that would normally require human intelligence. In the world of mining, AI is viewed as the next step in the digital transformation of mines, one that may affect everything from prospecting and exploration to the mining process itself (Schilling et al., 2017; Skilton and Hovsepian, 2018).

The potential benefits that AI may bring to different industries are immense, and include productivity increases and reduced costs (Purdy and Daugherty, 2016; Russell and Norving, 2015). Furthermore, the use of AI in the mining industry is nothing new: one of the first practical applications of AI in the 1970s was the development of expert decision-making systems that emulated the capacity of a human expert (Giarratano and Riley, 2004).

An expert system is designed to solve complex problems by reasoning using knowledge bases that initially stored and represented human knowledge via “if… then…”-type rules (rather than conventional computer programs). These are then processed by inference engines.

10%
THE EFFICIENCY SAVINGS THAT AI BRINGS

One of these early expert systems was PROSPECTOR, designed by the Stanford Research Institute (Hart, Duda, and Einaudi, 1978) for exploration in the mining sector. The system, which was initially conceived as a geological adviser, attempted to represent the knowledge and reasoning of expert geologists to automatically suggest possible drilling sites. The earliest results had an unprecedented impact. PROSPECTOR was able to predict the existence of unknown molybdenum deposits in Washington where no group of experts had detected any.

To achieve this, PROSPECTOR drew on over 2000 rules provided by experts which captured knowledge on geological formations and rock and mineral types. It then used computational methods based on probabilistic reasoning to generate conclusions. Figure 1 shows an extract from an original summary as an example of the type of results that PROSPECTOR can reach.

However, expert systems are just one of many types of AI applications. AI has traditionally been developed around problems that entail general intelligence: reasoning, problem-solving, knowledge, planning, learning, natural language, perception, social behavior, and the ability to handle objects.

In the realm of mining automation, four subareas of AI are generating impacts and significant benefits.

1. Machine learning: this entails the development of computational methods that improve performance at experience-based tasks and thus allow

FIGURE 1
EXTRACT FROM A MINERAL EXPLORATION DECISION MADE BY PROSPECTOR.

[PROSPECTOR IS READY TO SUMMARIZE ITS INTERPRETATION.]

ON A SCALE FROM -5 TO 5,
MY CERTAINTY IN (MVTD) IS NOW: .8995
[THE CERTAINTY OF .8995 WOULD BE INTERPRETED AS PERHAPS MILDLY ENCOURAGING.]

IN SUMMARY, THE FOLLOWING HAVE BEEN OBSERVED:
THE MAJOR FAVORABLE CONCLUSIONS WERE:
(* SUGGESTIVE EVIDENCE FOR MVTD) (5.0)
(* PROSPECTIVE ORE BODY LIES IN CARBONATE SEDIMENTS) (3.0)
(* RIGHT HYDROTHERMAL ALTERATION) (3.0)
(* PROSPECTIVE ORE BODY CONTAINS GALENA AND SPHALERITE) (2.0)

Source: Hart, Duda and Einaudi (1978).

the relationships between complex data to be understood so that decisions can be made automatically (Bishop, 2011). These types of decision can be simple (e.g., predicting what temperature a machinery component will reach under certain conditions) or complex (e.g., detecting when a crusher will break down). Through machine learning, an algorithm can detect complex patterns based on thousands of variables, even in the complex operating environments of mines (Shalev-Shwartz and Ben-David, 2014).

2. Autonomous robotics and perception: unlike in traditional robotics, where control of the robot is programmed in advance by human experts, an autonomous robot carries out tasks and makes decisions without human control or intervention (Siegwart, Nourbakhsh, and Scaramuzza, 2011). This is particularly appealing in the mining sector, where factors related to performance, cost, and safety tend to limit the tasks that can be performed by human operators. On the other hand, the environment in which autonomous robots work can be challenging, as they often contain many variables that cannot easily be predicted. In the mining context, a fully autonomous robot may be able to obtain information from the environment, work for long periods of time without human intervention (except for maintenance), and prevent humans from having to work in potentially hazardous circumstances.

Automatic planning: this concerns the development of strategies or action sequences (plans) in highly dynamic environments, which are usually carried out by an intelligent agent, autonomous robot, or driverless vehicle. Unlike traditional control and classification problems, the solutions to an activity planning problem are complex and need to

FIGURE 2
GLOBAL IMPACT OF AI

86% OF CLIENT INTERACTIONS WILL BE HANDLED WITHOUT HUMAN INTERVENTION BY 2020 (GARTNER).	US$1.5B HAD BEEN INVESTED BY FIRMS IN AI AS OF MAY 2016.	16% OF JOBS WILL BE REPLACED BY AI BY 2020 (FORRESTER).
1196 THE NUMBER OF AI STARTUPS IN THE WORLD (CRUNCHBASE).	6 B THE NUMBER OF ONLINE DEVICE THAT WILL MAKE SERVICE REQUESTS IN 2018 (GARTNER).	2 M THE NUMBER OF EMPLOYEES THAT WILL HAVE TO USE HEALTH AND SAFETY MONITORING DEVICES AS AN EMPLOYMENT CONDITION (GARTNER).

Source: Compiled by the author.

be reached and optimized in a scalable fashion so that a plan (action sequence) can then be generated and implemented to meet a new target (Ghallab, Nau, and Traverso, 2004). In unknown, dynamic environments, an automatic planning strategy would not only generate action plans but also review and update these automatically (e.g., the action plan for operators in different sections of a mine), so the process is iterative and seeks to reduce error margins and modify certain pre-established criteria.

Intelligent data analysis: this consists of analyzing large volumes of different types of data, dynamics, and information sources to detect or discover patterns or relationships within the data to convert them into a structure that can be understood by experts when making decisions (Keane, 2017). Intelligent data analysis usually uses machine and statistical learning methods and is closely connected to other fields such as data mining, business intelligence, and big data. Technologies for data analytics are widely used in commercial industries for making more informed decisions. The results of these types of analysis can help different industries to increase profits, improve the efficiency of their operations, optimize their marketing and customer service campaigns, respond faster to emerging market trends, and gain competitive intelligence.

The good news is that unlike in the early days of AI, several decades ago, these developments have become much more accessible and affordable for industries. AI can imply efficiency savings of up to 10% at some mining tasks, such as prediction, without requiring major injections of capital, but rather by simply producing better prediction and monitoring models (Walker, 2017) that improve efficiency as more and better data becomes available. The global impact that AI is having on different fields and

20%
REDUCTIONS
IN ACCIDENTS
AND TRANSPORTATION
COSTS

activities in the world has put it at the heart of the so-called Fourth Industrial Revolution. Figure 2 shows some indicators of what this technology means in innovation terms.

In the world of mining, it is estimated that improvements in some AI technologies such as data analytics and autonomous robotics could lead to annual savings of between US$290 billion and US$390 billion in annual savings on fuel, natural gas, thermal coal, iron ores, and copper by 2035 (Chui and McCarthy, 2018).

Generally speaking, there are four fundamental reasons for the technological transformation of mining using AI.

1. Real-time data harvesting. Data is obtained from high-precision sensors and instruments included in different types of machinery (e.g., drilling rigs). This speeds up the planning of multiple mining activities and includes intelligence in the decision-making process.

2. Creating an eco-friendly environment. Wireless systems and devices can be used to monitor environmental parameters to help assess the impact of different mining activities.

3. Reducing mining risk. Worker safety is improved by using remotely operated automatic drilling technology or autonomous robots. These AI methods can help operators and maintenance staff predict times when critical equipment will not be in use and anticipate potential indicator increases (e.g.,

pump pressure).

4. Simplifying mining operations. Intelligent autonomous robotic systems can carry out a wide variety of tasks (drilling, loading, explosions, and transportation), which can be supervised remotely and thus have a positive impact on the safety of human operators.

This technological transformation has enabled a fundamental shift in the way that the mining industry operates, one that is marked by the leveraging of information flows to reduce variability in decision-making and the use of more centralized mechanized operations to reduce variability in implementation. Table 1 shows examples of typical mining tasks and AI technologies that could be used to solve the problems that they entail.

The use of AI is becoming a cross-cutting feature of the entire mining cycle, from the exploration and extraction stages through to the end client, as is shown by the experiences of many mining companies around the world.

Multiple benefits have been observed in the industry, including the optimization of materials and equipment flows, improvements in advance troubleshooting or maintenance, increased mechanization through task automation, and real-time performance monitoring. This is why gradual applications of AI technologies in the global mining industry have been successful at solving decision-making problems and automating the different stages of the mining process.

MINERAL EXPLORATION

Exploration is a critical part of mining operations. A company could build the most the most advanced automated mining operation in the world, but this would be useless unless there is material in the ground to be extracted. The application of AI to prospecting and exploration in the mining sector is a very recent phenomenon and is attracting a lot of interest in the industry.

Some mining companies, such as Goldspot Discoveries, have developed

TABLE 1:
TYPICAL MINING TASKS AND THE USE OF AI IN RELATED PROBLEMS

TYPICAL MINING TASK	QUESTION	AI TECNOLOGIES	APPLICATIONS
PROSPECTING AND EXPLORATION	WHERE SHOULD WE EXPLORE?	MACHINE LEARNING	AUTOMATIC IDENTIFICATION OF ROCK AND SOIL TYPES, PREDICTIONS OF POTENCIAL LOCATION OF MINERAL ORE
DISCOVERY AND ADVANCED EXPLORATION	WHAT IS IN THE GROUND?	DECISION MAKING AND ANALYTICS	PREDICTION OF DRILLING SITES, CLASSIFICATION OF SUBSURFACE MINERALS
DEVELOPMENT AND CONSTRUCTION	HOW SHOULD WEBUILD THE MINE?	AUTOMATIC PLANNING	PRDICTION OF CONSTRUCTION STAGES USING MINERAL PROSPECTING
OPERATIONS AND PRODUCTION	HOW SHOULD WEREMOVE, MILL AND PROCESS MINERALS THAT ARE DISCOVERED?	ARTIFICIAL VISION, AUTONOMOUS ROBOTICS	TEXTURE BASED IMAGED SEGMENTATION, DRIVERLESS VEHICLES
LAND REHABILITATION	HOW CAN WE RESTORE THE SITE AND REPAIR DAMAGE TO THE ENVIRONMENT?	MACHINE LEARNING	PREDICT CHANGES IN CROSION PATTERNS, MONITOR ANIMAL MIGRATIONS, IDENTIFY ENVIRONMENTAL RISKS

Source: Compiled by the author.

machine learning algorithms and data mining tools to significantly improve their mineral exploration activities both locally and regionally. This investment decision model is used to acquire projects and royalties and invest in private funds to create an asset portfolio that is more profitable at a certain risk rate. The model can predict more than 80% of existing gold deposits in some regions using geological, topographic, and mineralogical data.

Furthermore, Kore Geosystems is planning to install instruments at drilling sites to provide real-time data that will help to speed up multiple stages of the mining process by providing decision-making intelligence. In fuel and gas, Yandex and Gazprom Neft are developing big data analytics methods by applying AI techniques during drilling and completion. To achieve this, the outcomes of different actions are estimated based on the data that is extracted from a given mine shaft to better prioritize the tasks and treatments that are carried out at the mine.

Large gold mining companies such as Goldcorp have teamed up with IBM and are using Watson, IBM's AI, to find better exploration sites in Ontario using vast quantities of geological data (Fatima, 2017). Petroleum and fuel companies have been using similar systems for years (Regulski, Szeliga, and Kusiak, 2014). Watson has enabled over 60 years of unstructured data from the Red Lake mine to be analyzed and interpreted to help geologists determine which areas might be valuable or to warn them of potentially unsafe situations. Watson will learn to think like a geologist and find patterns that had not been noticed before.

Another relevant aspect is explosion planning: measuring the impact of a high-energy explosion in a mine when each explosion is different and can never

13%
SAVINGS
ON BALLAST
CLEANING
COSTS

be repeated is a highly complex task. To explore this, companies have recently started using machine learning to understand the relationship between drilling patterns, explosion design, type of explosive, geology, and so on for more than 80 explosion events over the course of several months (Rosienkiewicz, Chlebus, and Detyna, 2017). The model was able to make predictions on the fragmentation that would have taken place using different types of explosives. This information was subsequently used to choose the right explosives so as to minimize the cost of the rock fragmentation that was being sought.

OPERATIONS AND PROCESSING

One of the most logical benefits of using AI in mining is improving the efficiency of operations. Many use the same basic advances in robotics and intelligent sensors that have been used in different factories to improve their performance. Mining, like any heavy industry, is an ideal place for the commercial use of driverless vehicles as they move slowly, operate in well-defined and controlled areas, and there is no need for concern about unauthorized personnel crossing the tracks. The Rio Tinto mining company has been at the forefront of the use of this type of technology and has expanded its fleet of around 80 autonomous trucks to its mining operations in Australia. The benefits to the company have been significant: these trucks are approximately 15% cheaper than human operators. Furthermore, they can operate around the clock without having to change shift (FN Media Group, 2017). However, they are not fully autonomous and rely on high-precision GPS for determining their locations and detect obstacles using laser sensors and radars.

BHP has recently replaced 50 of the trucks used at its mine in Jimblebar, Australia, with autonomous Caterpillar 793F models. The new fleet has reduced potential accidents and transportation costs by nearly 20%. This was made possible by the company's partnership with Caterpillar and the National Robotics Engineering Center at Carnegie Mellon University, Pittsburgh. The center helped design driverless freight vehicles for BHP and other mining companies such as the Fortescue Metals Group, for use in their iron mines (Tyler et al., 2017).

Unlike earlier experiences at the surface, Volvo began to trial completely autonomous underground transportation at its mine in Kristineberg, Sweden. This was a major challenge in comparison to the transportation used at other mines because underground vehicles cannot rely on help from GPS like vehicles on the surface can. Using GPS, a truck can navigate very narrow tunnels with great precision.

The Advanced Mining Technology Center (AMTC) in Chile is also making progress on driverless vehicles for use underground, on the surface, and in the air. The center mainly helps the mining industry to develop the automation, remote operation, and robotization of vehicles and mobile mining machinery. One of its more recent developments involved using drones to prospect for georesources (minerals, water, etc.), create topographic and magnetic mapping, model 3D slopes, and take environmental measurements, especially in topographically complex terrain such as inside underground mines and pits.

A further advantage of driverless land and air vehicles is that they are more predictable in how they perform such as hauling, which can lead to major savings in maintaining vehicles and transportation networks. This is because driverless vehicles use brakes and other controls more smoothly and predictably than humans do (Jiang et al., 2017).

Another use context is medium-sized smelting furnaces, where for years engineers have been trying to optimize products by applying their knowledge of chemistry and physics. However, the complex and shifting relationships between data have made this impossible up to now. Consequently, some mining companies have started using machine learning models to make predictions based on artificial neural networks that analyze vast quantities of data. It was observed that this type of model generates increasingly low results in the outcome of the chemical recipe, so one furnace implemented a new set of operational criteria that did not require investment but improved its outcomes by 2%.

DRILLING AND AUTONOMOUS TRANSPORTATION

Many mining companies not only use driverless vehicles but are also trying to make their entire operation autonomous. Some mining companies, such as Rio Tinto in Australia, use autonomous loaders, which collect waste, and autonomous drilling systems. The drilling system allows remote operators to control multiple drilling rigs. Recent tests show that these may improve productivity by nearly 10%.

Another common problem is that large quantities of material tend to be removed from the ground even when the mineral in question only accounts for a small proportion of this material (Ghasemi, Ataei, and Shahriar, 2014). Separating the sought-after material from useless waste, rocks, clay can become a very costly step in the mining process (Ibrahim, Bennett, and Campelo, 2015). To improve this, TOMRA has developed intelligent sorting machines that use machine learning technology. The system uses color, x-ray, and infrared sensors to examine any piece of material that moves through it and then sorts it based on criteria established by the company. The continual use of the system at some mines (e.g., Boliden) has resulted in nearly 12% less material needing to be removed. This means that less fuel and energy are consumed during processing and fewer haul trucks need to be loaded (Soofastaei et al., 2016). This builds on another recent achievement: the discovery of a 227-carat diamond at the Lula mine in Angola.

Chile's state mining company, Codelco, has also been using driverless haul vehicles that use precision GPS; artificial vision systems that take automatic, online granulometry measurements of material on conveyor belts and calculate froth flotation characteristics; and predictive model-based control, among other technologies.

Some mining companies in Africa have gone much further by using low-cost spatial monitors and drones to capture real-time information on truck locations, weather, speed, and vibration. They have used this data to design machine learning models based on artificial neural networks to analyze truck dispatching and monitor vehicle movements. The result is direct feedback that shows operators how they have been driving their trucks, which has helped them keep within the speed limit, cut down on short stops, and avoid sudden braking. This has reduced fuel consumption by 7%.

In other applications, freight train operators spend around 20% of their annual maintenance budget on ballast cleaning. To build a model that would be able to predict ballast deposits, concept trials were carried out using machine learning methods that drew on datasets obtained from radars, maintenance work, and weather monitors. This was used to design an optimization tool that helped identify the best sections of track to focus on. Trials showed a nearly 13% reduction in ballast cleaning costs and the elimination of almost any unnecessary maintenance.

EQUIPMENT AND MAINTENANCE MONITORING

Many equipment monitoring tasks in a mine have been made possible by cheap online sensors. Analyzing the data captured by the sensors using AI techniques can significantly improve maintenance, reduce downtime, and help predict problems before they happen.

Companies like General Electric and PETRA have developed this type of technology for various mining companies around the world. For example, PETRA's intelligent mining algorithms have allowed Newcrest Mining to significantly cut down on the number of overloading incidents at their semi-autogenous mills.

At other companies in Western Australia, AI has helped improve decision-making processes around machinery and supply chains. For example, some mining companies use expert systems to program track movements and dispatch the trains that haul iron between different mines. This application has significantly reduced cancellations due to congestion and means that more trains can operate at a given time (Wilk-Kolodziejczyk et al., 2017). These systems could perform even better by using advanced sensors and real-time process control, which would improve the quality and grade of minerals delivered to processing plants, while reducing water and energy usage. One example of this is the precision mining project carried out by the Escondida mining company in Chile. Through this initiative, the company is exploring ways of maximizing copper production and extending mine life using intelligent sensor technology in its heavy machinery, so as to analyze copper grades quickly and precisely. Recent tests used also this technology to produce smart caps that measured driver fatigue by analyzing their brainwaves for a fleet of more than 150 trucks.

With regard to maintenance tasks, mining equipment usually shows signs of wear and tear (pressure or temperature increases, electrical signals, noise, etc.) a long time before it breaks down. The use of low-cost sensors could thus allow large amounts of data on the state of equipment to be captured. As this could lead to information overload among executives and engineers, intelligent algorithms have been applied to automatically detect the unique signs of a potential equipment failure, modeling the relationship between the failures observed and data on factors that influence the state of equipment (driver behavior, historical maintenance, weather, etc.). In some cases, these methods can detect failures days in advance, which allows mining companies to program future maintenance work efficiently (Stefaniak, Wodecki, and Zimroz, 2016). This has led to increases in equipment usage time and a greater share of planned maintenance.

Despite the evident improvements in productivity and efficiency that AI can bring to specific mining tasks, the complete benefits of this technology can only be quantified when mining operations are integrated and managed as a sole system that runs from the mine to the market (Tyler et al., 2017). Furthermore, miners, engineers, and mining executives will have access to knowledge and information that are essential to making better decisions and, as a result, mine safety will improve as new ways are found to remove human operators from potentially dangerous activities or areas. Making operations more predictable can help companies reduce downtime, make better use of equipment, optimize production, and decrease stock shortages.

UNEVEN BENEFITS

Mining is all about producing commodities, so being competitive means producing products faster and cheaper. As the industry has focused on improving productivity and efficiency, it is no surprise that many mining companies have started to aggressively use AI to find ways of improving these efficiency levels. Some companies are already seeing tangible benefits from the use of smart machinery that is operated autonomously. This is also observed in the types of investment required: the marginal cost for adding an autonomous system to a haul truck is negligible.

While the application of AI has attracted the attention of the major stakeholders in the mining industry, innovation is still in its early stages and it is hard to envisage the scale of the impact it will eventually have. For progress with AI in heavy industry to be faster, companies may need to change the way that they go about research and development. For example, many mining companies are using the same hardware as each other and carrying out the same tests, so developing centralized data repositories could bring considerable savings in time and money. This would imply companies creating information consortia and using them to create computational methods that meet their individual requirements.

Furthermore, AI-enhanced robotic devices and autonomous transportation systems could carry out a range of tasks that involve practically the entire mining production chain, from exploration and hauling to delivery to the client. However, the benefits will not be distributed evenly throughout the industry: companies with large assets and vehicle fleets have a lot more potential for capturing large quantities of data, using AI more effectively, and using new technology as they become mature.

Finally, there is much controversy around the changes that AI may bring to the labor force through mine automation, as this would effectively eliminate jobs. However, the jobs that are most likely to be replaced are precisely those that are most dangerous, unhealthy, or monotonous for human operators to perform, so these changes should lead to the creation of better-quality jobs.

REFERENCES

Bishop, C. 2011. *Pattern Recognition and Machine Learning*. Cambridge: Springer.
Chui, M., and McCarthy, B. 2018. *An Executive's Guide to AI*. McKinsey & Company.
Craig, B. 2017. "Mining Data to Improve Safety and Cut Maintenance Costs." BHP. December 20.
Duddu, P. 2014. "Ten Technologies with the Power to Transform Mining." *Mining Technology.* April 10.
Durrant-Whyte, H., Geraghty, R., Pujol, F., et al. 2015. "How Digital Innovation Can Improve Mining Productivity." McKinsey & Company.
Fatima, S. 2017. "Goldcorp and IBM Canada Launch Watson Artificial Intelligence Project." *CIM Magazine.* April 25.
FN Media Group. 2017. "How Robots and Artificial Intelligence Will Transform Mining." *PR Newswire.* December 14.
Ghallab, M., Nau, D., and Traverso, P. 2004. *Automated Planning: Theory and Practice*. San Francisco: Morgan Kaufmann.
Ghasemi, E., Ataei, M., and Shahriar, K. 2014. "An Intelligent Approach to Predict Pillar Sizing in Designing Room and Pillar Coal Mines." *International Journal of Rock Mechanics and Mining Sciences* 65: 86–95.
Giarratano, J., and Riley, G. 2004. *Expert Systems: Principles and Programming*, *4th edition*. Boston, MA: Thomson Course Technology.
Habrat, M., and Lisowski, P. 2015. "Applying Image Mining to the Analysis of Rock Structure Images." *Geology, Geophysics and Environment* 41(1): 86–90.
Hart, P., Duda, R., and Einaudi, M. 1978. "PROSPECTOR: A Computer-Based Consultation System for Mineral Exploration." *Journal of the International Association for Mathematical Geology* 10 (5): 589–610.
Ibrahim, A.M., Bennett, B., and Campelo, C. 2015. "Predictive Expert Models for Mineral Potential Mapping." In: M. Khosrow-Pur, editor. *Encyclopedia of Information Science and Technology, 3rd edition*. Pennsylvania: IGI Global.
Jiang, Y., Zhixiong, L., Yang, G., et al. 2017. "Recent Progress on Smart Mining in China: Unmanned Electric Locomotive." *Advances in Mechanical Engineering* 9 (3): 1–10.
Karakaya, E. 2017. "Why General-Purpose Technologies Matter in Innovation Systems: The Case of Artificial Intelligence in the Mining and Metal Producing Industry of Sweden." Paper presented at the DRUID Society Conference. New York, US.
Keane, R. 2017. *Data Analytics: Master the Techniques for Data Science, Big Data and Data Analytics*. Middletown: CreateSpace Independent Publishing Platform.
Purdy, M., and Daugherty, P. 2016. *Why Artificial Intelligence is the Future of Growth*. Accenture.
PwC. 2017. *Mine 2017: Stop Think... Act*. PwC.
Rao, A., and Verweij, G. 2017. *Sizing the Prize: What's the Real Value of AI for Your Business and How Can You Capitalise?* PwC.
Regulski, K., Szeliga, D., and Kusiak, J. 2014. "Data Exploration Approach Versus Sensitivity Analysis for Optimization of Metal Forming Processes." *Key Engineering Materials* 611–612: 1390–1395.
Rosienkiewicz, M., Chlebus, E., and Detyna, J. 2017. "A Hybrid Spares Demand Forecasting Method Dedicated to Mining Industry." *Applied Mathematical Modelling* 49: 87–107.
Russell, S., and Norving, P. 2015. *Artificial Intelligence: A Modern Approach, 3rd edition*. Harlow: Pearson.
Schilling, D., King, J., Wood, R., et al. 2017. "Mining Value in AI." Boston Consulting Group. October 5.
Shalev-Shwartz, S., and Ben-David, S. 2014. *Understanding Machine Learning: From Theory to Algorithms*. New York: Cambridge University Press.
Siegwart, R., Nourbakhsh, R. I., and Scaramuzza, D. 2011. *Introduction to Autonomous Mobile Robots, 2nd edition*. Cambridge: MIT Press.
Skilton, M., and Hovsepian, F. 2018. *The 4th Industrial Revolution: Responding to the Impact of Artificial Intelligence on Business*. Cham, Switzerland: Palgrave Macmillan.
Soofastaei, A., Aminossadati, S. M., Kizil, M. S., et al. 2016. "Reducing Fuel Consumption of Haul Trucks in Surface Mines Using Artificial Intelligence Models." In: N. Aziz, N. Kininmonth, J. Nemcik, et al., editors. *Proceedings of the 2016 Coal Operators' Conference*. Wollongong: The University of Wollongong Printery.
Stefaniak, P., Wodecki, J., and Zimroz, R. 2016. "Maintenance Management of Mining Belt Conveyor System Based on Data Fusion and Advanced Analytics." *Advances in Technical Diagnostics* 10: 465–476.
Tyler, L., Busuttil, S., Rivieres, J., et al. 2017. "Value Creation Through Exploration in BHP." *AusIMM Bulletin.* June.
Walker, J. 2017. "*AI in Mining: Mineral Exploration, Autonomous Drills, and More*." *OperatorsTech Emergence Magazine.* December 3.
Wilk-Kolodziejczyk, D., Regulski, K., Gumienny, G., et al. 2017. "Data Mining Tools in Identifying the Components of the Microstructure of Compacted Graphite Iron Based on the Content of Alloying Elements." *The International Journal of Advanced Manufacturing Technology* 1–13.

Agricultural *Commodities*

Estimating Prices Using Artificial Intelligence

Julia M. Núñez Tabales, José M. Caridad y Ocerín, and María B. García-Moreno
University of Córdoba, Spain

ARTIFICIAL INTELLIGENCE–BASED MODELS FOR PREDICTING DAILY VARIATIONS IN COMMODITY PRICES ARE SIGNIFICANTLY MORE ACCURATE THAN TRADITIONAL MODELS. THIS ARTICLE ANALYSES EFFICIENCY PARAMETERS AROUND PRICE FORMATION FOR THE WHEAT AND CORN MARKETS.

Commodities are products that are in their raw state or have only been subject to minimal processing. Interest in analyzing commodity markets has surged in recent years (Belousova and Dorfleitner, 2012). From the producer to the final consumer, trade in commodities involves several stages in which a range of stakeholders are involved, including those that facilitate financing[1].

Trade in commodities is characterized by high levels of uncertainty around market prices. Consequently, research into the factors that determine peaks and troughs in commodity prices and the volatility of these (Deaton and Laroque, 1992) is essential to numerous groups that are involved in this market, notably producers themselves, investors, traders, and political players (Karali and Power, 2009). Accurately price prediction allows better decisions to be made regarding the right time to buy and sell and thus helps limit risks.

Commodity markets can be divided into five categories (London Stock Exchange, 2018): agricultural (wheat, corn, soy, among others); livestock (beef, pork, and so on); precious metals (gold, silver); industrial metals (aluminum, copper, and zinc); and energy (natural gas, oil, and electricity).

Taking the fundamentals of economic theory as a starting point, commodity prices can be analyzed for different purposes: setting future prices, analyzing their volatility, and validating market efficiency or value at risk. Other authors, in contrast, focus on examining the internal statistical behavior of price series (Coronado Ramírez, Ramírez Grajeda, and Celso Arellano, 2012).

If we assume the efficient market hypothesis to be true for these products, it would be impossible to predict variations in prices based on the past. If the time series of prices for a commodity is taken as yt , the weak form of the efficient market hypothesis establishes that the only model for representing this series would be a random walk, in other words, $y_t = y_{t-1} + a_t$, where the stochastic process at is a series of independent, homoscedastic, and nonautocorrelated random variables. One abbreviated way of representing this process is by using a difference operator, the result of which is $\nabla y_t = a_t$. The efficient market hypothesis (Samuelson, 1965) argues that, given a set of information I(t) available at point t, the prices, y_t, of a product verify that the expected value of changes in these prices is nil:

$$E[y_{t+1} - y_t \mid I_t] = E[\nabla P_{t+1} \mid I(t)] = 0$$

so the best prediction that can be made for the next day's price is today's price.

However, in this article, we test different modeling hypotheses for wheat and corn that may be useful when market inefficiencies have been detected. The markets for these products should, in principle, behave efficiently, so different stakeholders should not be able to anticipate changes in prices. Despite this, markets provide different parties with asymmetrical information. However, everyone has access to past prices, so weak efficiency can be analyzed. A significant number of studies attempt to contrast weak efficiency in stock markets. Some recent publications in this direction have included Mynhardt, Makarenko, and Plastun (2017), Serin (2017), and Tang et al. (2017). In Europe, similar studies include Dicle and Levendis (2011), Caraiani (2012), Khan and Vieito (2012), Apergis, Artikis, and Kyriazis (2015), and Gupta and Sankalp (2017). South American markets have been the focus of studies such as that of Ojeda Echeverri and CastYEA Vélez (2014). The Asian market is analyzed in works such as Hamid et al. (2010), Lean, Mishra, and Smyth (2015), Shaik and Maheswaran (2017), Hou et al. (2017), Soon and Abdul-Rahim (2017), and Gupta and Singla (2018). Likewise, African markets have been studied by authors such as Mazviona and Nyangara (2013) and Ikeora, Charles-Anyaogu, and Andabai (2016).

TRADITIONAL STATISTICAL MODELS

Predicting commodity market prices and production first began to be studied more than a century ago. The first econometric prediction model for agricultural commodities was presented in 1917 (Allen, 1994). Since then, in the sphere of agricultural economics and finances, numerous econometric models for predicting commodity prices based on time series have been put forward. It is worth noting the use of the Box-Jenkins methodology (Kohzadi et al., 1996; Ntungo and Boyd, 1998), which entails autoregressive (AR), moving average (MA), and autoregressive integrated moving average (ARIMA) models, and to the use of vector autoregression (VAR) (Alonso and Arcila, 2012), transfer functions and dynamic analysis (Aradhyula and Holt, 1988), generalized autoregressive heteroscedastic (GARCH) models (Adrangi and Chatrath, 2003; Villada, Cadavid, and Molina, 2008; Chang, McAleer, and Tansuchat, 2009; Benavides Perales, 2009), and smooth transition vector error correction models (STVECM) (Milas and Otero, 2002).

Some publications have demonstrated the chaotic behavior of prices using different techniques, such as the Lyapunov exponent test (Cromwell and Labys, 1993), the BDS (Brock, Dechert and Scheinkman) statistic test (Ahti, 2009), the correlation exponent (Tejeda and Goodwin, 2009), or even artificial neural networks (ANN) (Velásquez Henao and Aldana Dumar, 2007).

Combined or hybrid models reached by aggregating two or more models are increasing predictive capacity and have recently begun to predominate in the literature. These models allow for the use of different structural characteristics (such as qualitative variables), take nonlinear features into account to describe the abolition of prices, and are characterized by their flexibility.

Ribeiro and Oliveira (2011) put forward a model that combines ANNs with stochastic methods (the Kalman filter) to predict sugar prices on the Brazilian and Indian markets. Zou et al. (2007) propose a hybrid model using ANNs and ARIMA to predict wheat prices on the Chinese market and reach the conclusion that the best fit is achieved by the network on its own. Zhang (2004) and Tseng et al. (2008) used the same combination of models. In contrast, Sallehuddin et al. (2007) used a complex Grey relational artificial neural network (GRANN) model in combination with ARIMA to analyze the evolution of crop yields in China. More recently, Kristjanpoller and Minutolo (2015) put forward a hybrid GARCH-ANN model to predict the volatility of gold prices.

These hybrid models can be divided into three categories: conventional, hybrid I, and hybrid II (see table 1) (Sallehuddin et al., 2007; Ruiz-Gándara and Caridad y Ocerín, 2014). Conventional and hybrid I models use the same sequence of hybridization that is applied to a linear model to find linear relationships between data. In what follows, we use ANNs to try to model the residuals derived from a linear model. In this case, the assumption is that the linear components have been fully identified by the linear model and that the residuals thus contain the nonlinear component.

Likewise, the hybrid II and the conventional hybrid models follow a sequence of inverse hybridization. In the hybrid II model, GRANN is applied initially, followed by a linear ARIMA model. In this model, GRA is used to select significant factors before predictions are made using ANN.

ARTIFICIAL NEURAL NETWORKS

Artificial intelligence attempts to imitate the intelligent behavior usually associated with human beings and encompasses techniques as varied as fuzzy logic, expert systems, genetic algorithms, or ANNs.

Over the last two decades, ANNs have been used for a variety of applications in connection with different fields of study (Jain and Kumar, 2007). Various ANN applications have been used to solve prediction problems for complex nonlinear time series, such as demand for electricity (Abraham and Nath, 2001), electricity prices (Ganeta, Romeo, and Gil, 2006), stock market index volatility (Hamid and Iqbal, 2004), or rainfall (Srinivasulu and Jain, 2006). ANNs have also been

TABLE 1:
FEATURES OF THE DIFFERENT TYPES OF HYBRID MODELS

TYPE OF HYBRID	DATA	MODEL	HYBRIDIZATION SEQUENCE	SELECTED FEATURES
Conventional	Univariate	ARIMA, RNA	Linear, nonlinear	None
Hybrid I	Multivariate	MR, RNA	Linear, nonlinear	Goodness-of-fit test
Hybrid II	Multivariate	GRANN, ARIMA	Nonlinear, linear	Grey Relational Analysis (GRA)

Source: Sallehuddin *et al.* (2007).

shown to be effective for making predictions using both noisy and noise-free time series (Zhang, Patuwo, and Hu, 2001). Hippert, Pedreira, and Souza (2001) and Zhang (2004) have stated that prediction is undoubtedly one of the main areas of application for ANNs.

ANNs are nonlinear models that consist of a structure made up of nodes or neurons. These neurons are structured into layers and are connected by weights that determine the intensity of the connections. These weights are the parameters for the model and are obtained through optimization techniques that minimize certain measures of error.

The most frequently used ANN models include the so-called feed-forward model, the neurons in which that are not connected to those in the previous layer or same layer—in other words, there is no feedback cycle. The network used most for economics and finances is the multilayer perceptron (MLP) network, which is made up of an input layer, one or more hidden layers, and an output layer.

Building a neural network entails selecting an architecture, that is, establishing the number of layers, the number of neurons within each layer, and the activation functions. The most frequently used activation functions include the sigmoid and the hyperbolic tangent (Villada et al., 2008). The training algorithm also needs to be chosen, the most common of which is the backpropagation algorithm. Funahashi (1989) and Hornik, Stinchcombe, and White (1990) demonstrated that a network that might be considered a universal approximation function can be obtained using only a hidden layer and a sigmoid activation function. Furthermore, as Tu (1996) argues, there is no theoretical consensus that predetermines the number of neurons in the hidden layer, so determining this entails a process of trial and error.

TABLE 2:
STUDIES USING ANNS ON COMMODITY SERIES

AUTHOR	YEAR	COMMODITY
Kohzadi *et al.*	1996	Wheat and livestock
Ntungo and Boyd	1998	Corn, silver, Deutsche Mark
Yonenaga and Figueiredo	1999	Soy
Freiman and Pamplona	2005	Beef
Zou *et al.*	2007	Wheat
Velásquez Henao and Aldana Dumar	2007	Coffee
Villada, Cadavid and Molina	2008	Electricity
Liu	2009	Gold
Malliaris and Malliaris	2009	Gold, petroleum, and Euros
Yu and Ou	2009	Tomatoes
Ferreira *et al.*	2011	Soy, beef, corn, wheat
Ribeiro and Oliveira	2011	Sugar
Pinto	2012	Corn, petroleum, gold, and copper
Wiles and Enke	2014	Soy
Kristjanpolle and Minutolo	2015	Gold
Villada, Muñoz and García-Quintero	2016	Gold
Pinheiro and Senna	2017	Sugar and soy

Note: Studies on agricultural commodities are marked in blue.
Source: Compiled by the authors.

FIGURE 1:
RANDOMNESS TEST SEQUENCES

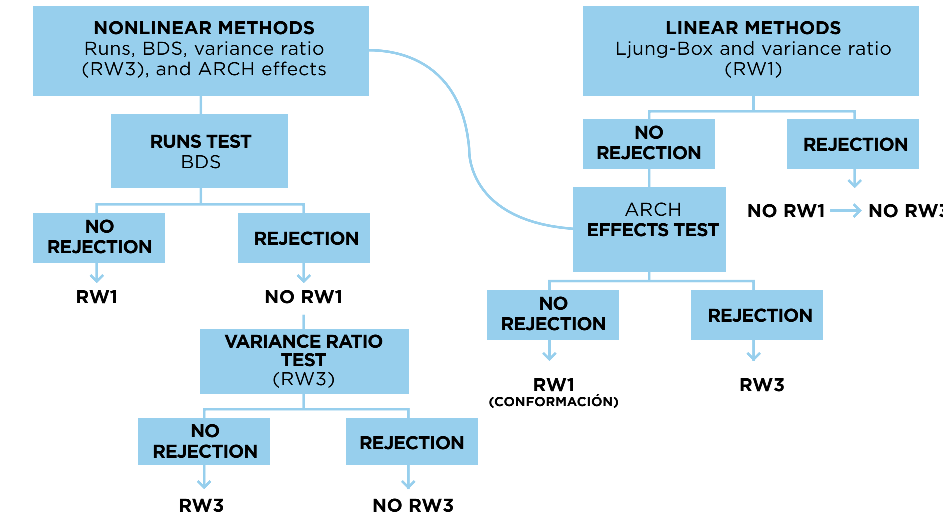

Source: Compiled by the authors.

Although most of the research published demonstrates the superiority of ANN models in comparison with traditional linear models, various studies reach different conclusions regarding the performance of ANNs (Sallehuddin et al., 2007). For example, Denton (1995) and Gorr, Nagin, and Szczypula (1994) showed that ANNs achieve a very similar degree of fit to that of the linear model. On the other hand, Brace, Schmidt, and Hadlin (1991), Caire, Hatabian, and Muller (1992), Heravi, Osborn, and Birchenhall (2004) and Taskaya-Temizel and Casey (2005) concluded that ANNs were not as effective at predicting as linear models. However, Kang (1991) showed that ANNs always function better than ARIMA models and their performance improves even more the longer the prediction horizon is.

Table 2 summarizes some of the major contributions of the literature that uses ANNs to predict commodity prices. Studies on agricultural commodities are marked in blue.

METHODOLOGY

This study analyzes various agricultural commodities through daily changes in prices over an interval of more than two years, in order to contrast market efficiency and use different explanatory time-based models to attempt to predict variations in prices based on the past and the evolution of prices of other products. The models used are multiple equation models and are based on linear techniques, such as VAR models, and nonlinear techniques, such as ANN models. We used a VAR model to estimate variations in the price of various products based on variations in the immediate past. For example, taking the daily prices for two commodities x_t and y_t, the variations in these compared to the previous day ∇x_t and ∇y_t are related to the variations in prices during the p previous days:

$$\nabla x_t = a_{x1}\nabla x_{t-1} + a_{y1}\nabla y_{t-1} + a_{x2}\nabla x_{t-2} + a_{y2}\nabla y_{t-2} + \dots + a_{xp}\nabla x_{t-p} + a_{yp}\nabla y_{t-p} + a_t^*$$

$$\nabla y_t = b_{x1}\nabla x_{t-1} + b_{y1}\nabla y_{t-1} + b_{x2}\nabla x_{t-2} + b_{y2}\nabla y_{t-2} + \dots + b_{xp}\nabla x_{t-p} + b_{yp}\nabla y_{t-p} + b_t^*$$

where the vector (a_t^*, b_t^*) is a bivariate white noise, that is, a centered nonautocorrelated random variable with a constant covariance matrix $\sum$ for all t. Other exogenous explanatory variables may be introduced into the model. The Granger causality test is based on this type of model (Alonso and Arcila, 2012).

Prior to the modeling process, a series of statistical contrasts must be implemented to decide if a series of prices can be said to have been generated by a random walk or by a martingale or if it shows a certain degree of time dependence. We thus distinguish between different processes to represent the price series: RW1, or random walk 1, in which the expression $\nabla y_t = a_t$, $a_t \in$ IID(0, σ^2) is a succession of independent and identically distributed (IID) random centered variables; RW2, or heterogeneous random walk, in which at are independent variables that are not necessarily equally distributed and thus may be centered and heteroscedastic; and RW3, in which at are uncorrelated centered variables. Logically, the first of these categories is the most restrictive, making it impossible to predict prices based on the past, while the last is the least restrictive.

A series of tests were applied to distinguish between these models. One initial step was to use the Ljung-Box autocorrelation test (Caridad y Ocerín and Caridad y López del Río, 2014) although in this case what we were trying to detect were linear relationships. Barnett et al. (1997), however, examine nonlinearity in many financial markets. The classic Wald-Wolfowitz runs test for randomness could also have been used as could the Lo and MacKinlay (1988) variance ratio test, in which variance in the series is compared over different periods. We needed to use nonlinear methods to complete the efficiency tests. Engle (1982) proposes a test based on second-order moments, considering ARCH(p)-type dynamic heteroscedasticity. Another option is the Brock, Dechert, and Scheinkman (1987) test, which is based on the correlation integral method developed by Grassberger and Procaccia. To distinguish between the RW1 random walk process and the RW3 correlation process, the series of tests shown in figure 1 were used. If the null hypothesis is accepted, the runs and BDS tests take the process as being RW1, although the rejection of these does not necessarily imply that the process is RW3. If the contrast hypothesis is accepted, then the variance ratio test, which is applicable when there is heteroscedasticity, is accepted as being RW3, although a rejection of the contrast hypothesis does not imply that it is necessarily not of this type. In short, the possibility of nonlinear relationships must be contemplated.

DATA ANALYSIS

This study uses data on the following commodities: corn (CA) and wheat (WA). Daily information is available on the former from December 2013 onward (a total of 1,565 entries), while the second series includes 1,191 entries from July 2015 onward. Both products have experienced downward trends in recent years, although the price of wheat has been the more volatile of the two. The correlation between the two series is high, equal to 0.856, although they do not present significant cross-correlations (figure 2).

The programs we used for the analysis were QMS EViews 10 and IBM's SPSS Statistics 23. The former was used to run different efficiency tests and VAR modeling and the latter

FIGURE 2
EVOLUTION OF DAILY PRICES

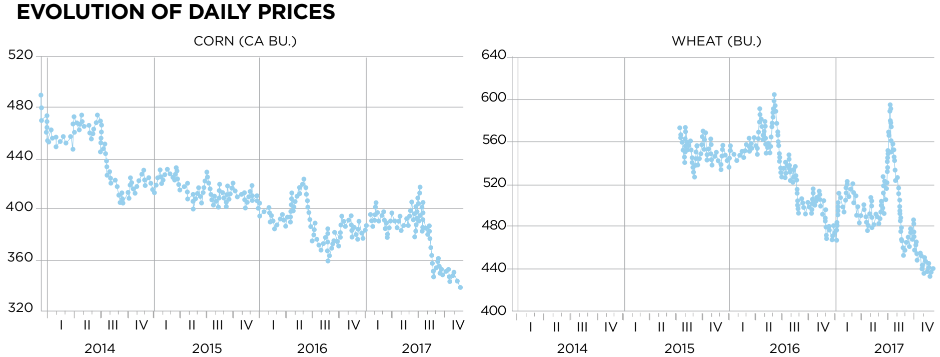

Source: Compiled by the authors.

TABLE 3
APPLICATION OF THE LJUNG-BOX TEST TO VARIATIONS IN PRICES

CORN CA							WHEAT WA						
AUTOCORRELATION	PARTIAL AUTOCORRELATION		r	r*	Q	p	AUTOCORRELATION	PARTIAL AUTOCORRELATION		r	r*	Q	p
		1	-0.045	-0.045	1.4842	0.223			1	0.073	0.073	3.8681	0.049
		2	-0.007	-0.009	1.5174	0.468			2	-0.044	-0.050	5.2989	0.071
		3	-0.100	-0.101	8.7920	0.032			3	-0.032	-0.025	6.0586	0.109
		4	-0.003	-0.012	8.7981	0.066			4	0.018	0.020	6.2960	0.178
		5	0.031	0.028	9.4873	0.091			5	0.059	0.054	8.8485	0.115
		6	-0.059	-0.067	12.011	0.062			6	0.029	0.022	9.4653	0.149
		7	-0.030	-0.037	12.661	0.081			7	-0.013	-0.011	9.5859	0.213
		8	-0.095	-0.095	19.240	0.014			8	-0.041	-0.034	10.802	0.213
		9	0.005	-0.018	19.260	0.023			9	-0.034	-0.030	11.635	0.235
		10	-0.005	-0.017	19.278	0.037			10	0.011	0.008	11.728	0.304
		11	-0.012	-0.032	19.390	0.054			11	0.058	0.050	14.178	0.223
		12	-0.011	-0.020	19.485	0.077			12	0.015	0.009	14.353	0.279
		13	0.028	0.024	20.077	0.093			13	-0.013	-0.004	14.477	0.341
		14	0.006	-0.010	20.107	0.127			14	-0.050	-0.042	16.354	0.292

Source: Compiled by the authors.

TABLE 4
APPLICATION OF THE ARCH TEST TO VARIATIONS IN PRICES

q	1	2	3	4
∇CA CORN	63.15003 (0.0000)	63.05135 (0.0000)	67.24453 (0.0000)	69.44998 (0.0000)
∇WA WHEAT	13.77034 (0.0002)	14.41457 (0.0007)	96.80639 (0.0000)	97.00701 (0.0000)

Source: Compiled by the authors.

to estimate ANNs.

EFFICIENT MARKETS?

Corn and wheat are, along with soy, undoubtedly the largest agricultural commodity markets. To analyze their efficiency, we applied different statistical contrasts to the series of daily prices in the following sequence: the Ljung-Box autocorrelation contrast, the dynamic heteroscedasticity test, the randomness runs test, the BDS test, and the variance ratio test.

Some contrasts generated low limit probabilities (p), although this may be due to the fact that the sample sizes were very large, which can sometimes lead to the alternative hypothesis being accepted erroneously.

The Ljung-Box test shows autocorrelation for variations in corn prices (∇CA_t) based on four-day lags, which begin to disappear after two weeks, although this may be partly due to the sample size, as suggested above. No autocorrelation was detected in the case of wheat (∇WA_t) except in the first two days and even this was less marked. This points to accepting the RW1 behavior hypothesis for wheat, although the situation is not so clear for corn (see table 3).

The ARCH(q) dynamic heteroscedasticity contrasts are shown in table 4.

Dynamic heteroscedasticity was detected, as was to be expected after observing the figures for ∇CA_t and de ∇WA_t, which show that the distribution of price variations changes over time (figure 3).

The runs test provides a limit probability p = 0.0048 for the series of variations in corn prices, which implies rejecting the RW1 behavior hypothesis. The wheat market, however, appears to be more efficient, as p = 0.1742.

FIGURE 3
VARIATIONS IN PRICES

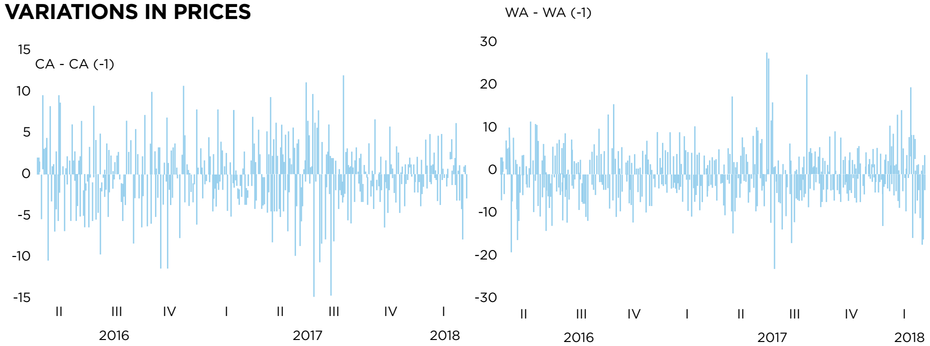

Source: Compiled by the authors.

TABLE 5
RESULTS FOR THE APPLICATION OF THE BDS TEST TO VARIATIONS IN PRICES

	m	BDS	z	p
∇CA CORN	2	0.014538	4.813934	0.0000
	3	0.008968	2.683526	0.0073
	4	0.004985	1.794235	0.0728
	5	0.002848	1.406389	0.1596
∇WA WHEAT	2	0.010429	2.727861	0.0064
	3	0.005301	1.274825	0.2024
	4	0.000311	0.091599	0.9270
	5	-0.001388	-0.570155	0.5686

Source: Compiled by the authors.

Table 5 shows the results of the BDS test for variations in the prices of both products.

Again, the series of corn prices has low limit probabilities for m = 2 and 3 and even for m = 4, which tends to reject the RW1 structure. The results for the wheat series, in contrast, from m = 2 onward, clearly point to accepting this efficiency structure, even though a low p value was obtained for m = 1.

For the variance ratio test (table 6), the results obtained for different values of k were along the same lines as above: the corn market shows clear signs of inefficiency, which suggests accepting that it diverges from RW3, while RW3-type evolution is accepted for the wheat market.

VAR MODELING

As there are indications of slight inefficiencies in the corn market, we attempted to predict variations in the prices of this based on past performance and variations in wheat market prices. As the two series are related, two types of multiple equation models were estimated: a VAR(p) model for different values of p and another using an ANN. The dependent variables are price variations at point t, and the explanatory variables are these same variations using p lags.

Logically, an increase in p improves the fit, although this does not imply

TABLE 6
APPLICATION OF THE VARIANCE RATIO TEST TO VARIATIONS IN PRICES

	k	VR(k)	z	p	GLOBAL
∇CA CORN	2	0.935923	-1.630904	0.1029	2,940210 (0.0131)
	4	0.845428	-2.288122	0.0221	
	8	0.743417	-2.662887	0.0077	
	16	0.601020	-2.940210	0.0033	
∇WA WHEAT	2	1.009323	0.200167	0.8414	0.652180 (0.9443)
	4	0.946077	-0.652180	0.5143	
	8	0.951421	-0.369821	0.7115	
	16	0.953029	-0.248255	0.8039	

Source: Compiled by the authors.

TABLE 7
MEASURES OF GOODNESS OF FIT FOR VAR MODELS

p=42	CORN	WHEAT	p=7	CORN	WHEAT
R-squared	0.145164	0.129583	R-squared	0.031815	0.019054
Adj. R-squared	0.032791	0.015162	Adj. R-squared	0.012697	-0.000316
Sum sq resids	6034.350	15829.87	Sum sq resids	6834.492	17840.00
SE. equation	3.073016	4.977236	SE. equation	3.104774	5.016194
F-statistic	1.291809	1.132510	F-statistic	1.664132	0.983707
Log likelihood	-1794.908	-2144.032	Log likelihood	-1839.982	-2187.307
Akaike AJC	5.193116	6.157546	Akaike AJC	5.124260	6.083721
Schwarz SC	5.731385	6.695816	Schwarz SC	5.219248	6.178709
Mean dependent	-0.036948	-0.144682	Mean dependent	-0.036948	-0.144682
S.D. dependent	3.124674	5.015403	S.D. dependent	3.124674	5.015403
Determinant resid covariance (dof adj.)	131.2808		Determinant resid covariance (dof adj.)	135.9529	
Determinant resid covariance	102.2647		Determinant resid covariance	130.3779	
Log likelihood	-3729.801		Log likelihood	-3817.721	
Akaike information criterion	10.77293		Akaike information criterion	10.62906	
Schwarz criterion	11.84947		Schwarz criterion	10.81904	
Number of coefficients	170		Number of coefficients	30	

Source: Compiled by the authors.

that the explanatory variables in question have predictive power (see table 7). Two alternative VAR models and ANN equivalents will be used, with p = 42 lags—in other words, the previous six weeks—and p = 7 lags. Different measures of goodness of fit will be used for each model which show how limited their predictive power is for the variations in the prices of the two products.

Even using a large number of lags, the model's coefficient of determination is 0.145 for corn and 0.13 for wheat; in other words, both equations have very limited predictive power. With seven lags, these coefficients are reduced to 0.032 and 0.019, respectively, which are insignificant. However, if they are compared with measures based on information criteria, such as the Akaike and Schwarz criteria, which penalize the overparameterization of the model, the equations estimated using seven lags are preferable for both products. In short, the coefficients for additional data over the last five weeks do not provide information that is of use for making (linear) predictions of variations in prices for the two products. Furthermore, the goodness of fit for the models using p = 7 was negligible, thus confirming the hypothesis that there is no correlation between the past and present prices of these products, nor can variations in price be explained by examining the variations in price for the other commodity. It is also worth introducing other explanatory variables into these models, such as variations in the prices of other commodities. Doing so substantially increases the coefficient of determination for the different equations, although these relationships are not thought to be causal and therefore cannot be used for prediction purposes.

TABLE 8
NEURAL NETWORK FIT AND RELATIVE ERRORS (P = 7)

TRAINING	Sum of squared errors		568.207
	Global mean relative error		.786
	Relative error for	CA	.792
	scale dependents	WA	.780
	Stopping rule used		relative change in error 0.0001 0:00:00,19
	Training time		

Source: Compiled by the authors.

ANN MODELS

The neural network models we used are similar to the VAR models described above in terms of the explanatory variables: price variations are introduced as endogenous variables at point t to be explained by the lagged values up to horizon p. Two models were tested, one using p = 43 days and another using p = 7.

The specification used networks with a hidden layer and a low number of neurons to avoid overparameterization. The activation function we used was the hyperbolic tangent, with a standardized scale for output variables and a scaled conjugated gradient algorithm. The results for each network are based on all the avail-

TABLE 9
ESTIMATIONS OF THE PARAMETERS FOR THE NEURAL NETWORK

		FORECAST					
		HIDDEN LAYER 1				OUTPUTAYER	
PREDICTOR		H (1:1)	H (1:2)	H (1:3)	H (1:4)	ZCA	ZWA
INPUT LAYER	(BIAS)	-.312	-1.302	.326	.953		
	ZCA1	-.916	1.078	1.564	.276		
	ZWA1	2.596	-4.729	-4.155	-.735		
	ZCA2	1.434	-.429	-2.550	-1.095		
	ZWA2	-.747	1.517	1.415	-.776		
	ZCA3	1.352	-1.393	-2.157	-1.614		
	ZWA3	-.754	.970	.819	2,181		
	ZCA4	-.831	-1.298	1.528	.817		
	ZWA4	.588	.817	-1.202	-.196		
	ZCA5	.913	-.716	-1.532	-.425		
	ZWA5	1.450	1.222	-2.145	-2.880		
	ZCA6	-3.071	.557	4.561	2.741		
	ZWA6	.768	-.152	-1.016	-.022		
	ZCA7	-.142	.793	.567	-1.898		
	ZWA7	-1.537	-.575	2.172	2.404		
HIDDEN LAYER	(BIAS)					.108	.088
	H(1:1)					3.469	3.460
	H(1:2)					.291	.302
	H(1:3)					2.868	2.747
	H(1:4)					.449	.620

Source: Compiled by the authors.

FIGURE 4
VALUES OBSERVED AND ESTIMATED FOR VARIATIONS IN PRICES (P = 42)

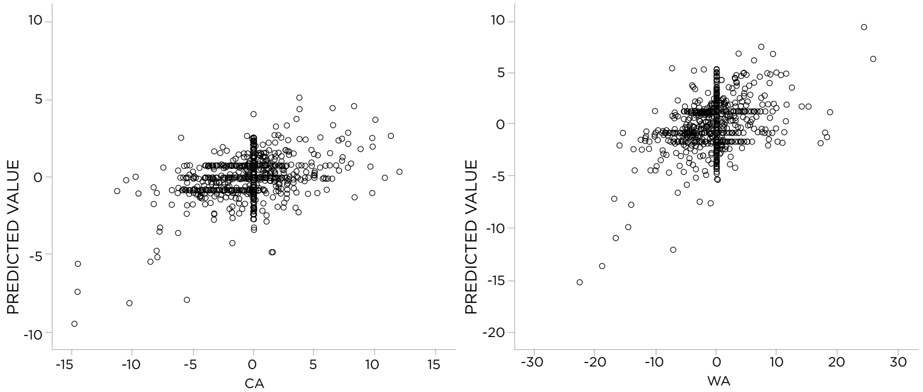

Source: Compiled by the authors.

able data so that they can be compared with the VAR models, although the results for predictive power were reached by re-estimating these same networks using 70% of the data as a training set. The parameters for the stopping rule are as follows: 0.0001 for the relative minimum change in the training error and 0.001 for the rate of these errors.

The network topology that we eventually selected (with p = 7 lags) has four neurons in the single hidden layer, as shown in the scheme below, in which the input variables ∇CA_{t-i} and ∇WA_{t-i} are displayed as ZCAi and ZWAi and shown for i = 1,2,..., 7, and the variables estimated for ∇CA_t and ∇WA_t. are ZCA and ZWA. Consequently, this network, like the VAR models above, was used to predict variations in prices for the following day.

The goodness-of-fit results are summarized in table 8, which shows that the relative error in the estimation is less than 0.8. However, when additional data was used to test the network, the predictive power decreased, and the relative predictive error was 0.97 and 0.971, respectively, for the two variables.

TABLE 10
FIT AND RELATIVE ERRORS IN THE NEURAL NETWORK (P = 42)

TRAINING	Sum of squared errors		283.581
	Global mean relative error		.404
	Relative error for scale dependents	CA	.388
		WA	.420
	Stopping rule used		relative change in error 0.0001
	Training time		0:00:00,59

Source: Compiled by the authors.

FIGURE 5
RELATIVE SIGNIFICANCE OF THE EXPLANATORY VARIABLES FOR ANN MODELS

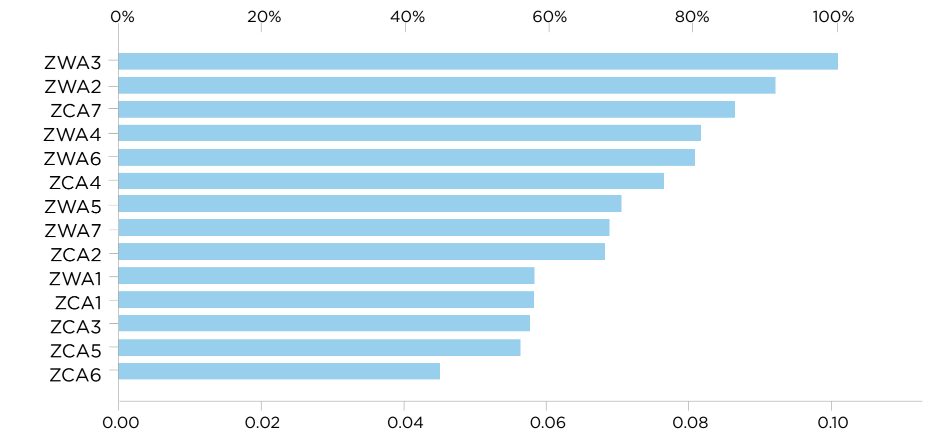

Source: Compiled by the authors.

The coefficients of determination for the two equations are 0.208 for corn and 0.22 for wheat, values that are low despite being significantly higher than those from the linear models.

The estimations of the parameters for the neural network are shown in table 9.

Figure 4 shows the values observed and estimated by the network. The network-based models have slight predictive power for prices, which is not observed in the VAR models. The significance of the explanatory variables that were included is shown in figure 5.

Based on this relative significance, there would seem to be a delay of two to three days and even as much as a week in the processing of information by the agents that play a part in the process.

The topology of the network that we ultimately selected (with p = 42 lags) contains four neurons in a single hidden layer (table 10).

However, if the network is used to make predictions with the set of data excluded from this estimation, the relative prediction errors increase to 0.867 and 0.881.

The coefficients of determination for the two equations are 0.612 and 0.58, respectively.

When the number of lags in the sets of explanatory variables increases, the predictive capacity of the ANN model we have proposed improves (in comparison with the network using p = 7), which explains a considerable part of the variance in price increases from one day to the next (figure 6). However, this predictive power suffers when these models are applied to

300%
THE IMPROVEMENT IN PREDICTIONS USING NEURAL NETWORK MODELS

FIGURE 6
VALUES OBSERVED AND ESTIMATED FOR VARIATIONS IN PRICES (p = 42)

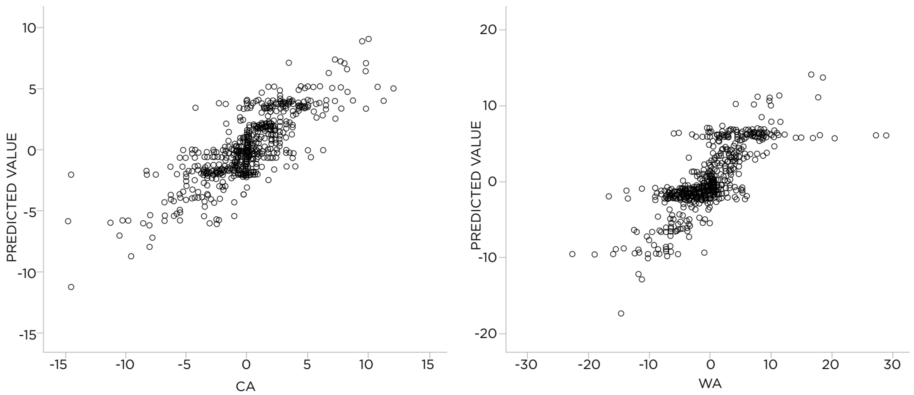

Source: Compiled by the authors.

data not included in the training set, which limits their applicability.

IMPROVING OUR PREDICTIONS

Agricultural commodity markets should behave efficiently, a priori, and the agents involved in them should thus not be able to anticipate future variations in prices. Even with a prediction horizon of one day, as was contemplated in this study, it should not be possible to use models to attain this objective. To consider this problem, we followed two approaches. The first entailed the application of statistical tests to different forms of random walks as a theoretical model that prices should follow. The second uses linear and nonlinear prediction models.

The first approach has been examined in the literature, fundamentally in stock markets, although the approach that we have taken is based on performing a series of tests to discriminate between various possible types of stochastic process. The second approach is based on the hypothesis that, if markets are not efficient, a model could be estimated to tackle the prediction problem.

After applying the proposed sequence of statistical tests to the series of variations in daily corn prices, we concluded that linear relationships existed that were incompatible with any of the forms of random walk that would imply the existence of efficiency in this market, and, furthermore, that the series is affected by dynamic heteroscedasticity. Even when the BDS test was applied to this series, market efficiency was rejected in the most restrictive way possible for RW1. However, the variance ratio test was not so conclusive as to allow us to reject RW3, which is the laxest form in which a market can be considered to be efficient.

In the case of the wheat market, there are no autocorrelations in price increases, a necessary but insufficient condition for any forms of RW. Consequently, after applying the ARCH effects test to detect dynamic heteroscedasticity, market efficiency in the form established by RW3 can be said to exist. Furthermore, the application of the BDS test to the series of increases in wheat prices allowed us to deduce that this is IID, so market efficiency in the form established by RW1 can be accepted. However, when the variance ratio test was applied to the same series, it led us to the conclusion that we could only accept the efficiency of the wheat market in the form established by RW3. In any case, the wheat market can be seen to be more efficient than the corn market, for which some price forecasting may be possible.

After estimating multiple-equation VAR-style models, our basic conclusion is that it is not possible to make approximate predictions regarding variations in the prices of the commodities in question. However, when nonlinear models like ANNs were used, relationships emerged that gradually improved the goodness of fit as the number of lags in explanatory variables increased. However, although we were able to model these interrelationships, the prediction problems cannot be solved without allowing for some highly significant relative errors. By introducing longer series and including blocks of several other agricultural products, the partial inefficiency of the markets for some of these products may allow us to solve these prediction problems more effectively.

NOTES

[1] The authors are part of the Department of Statistics, Operative Investigation, Econometrics, Business Management, and Applied Economics (DEEIOOEEA) at the University of Córdoba, Spain.

REFERENCES

Abraham A., and Nath B. 2001. "A Neuro-Fuzzy Approach for Modelling Electricity Demand in Victoria." *Applied Soft Computing Journal* 1 (2): 127–138.

Adrangi, B., and Chatrath, A. 2003. "Non-Linear Dynamics in Futures Prices: Evidence from the Coffee, Sugar and Cocoa Exchange." *Applied Financial Economics* 13 (4): 245–256.

Ahti, V. 2009. "Forecasting Commodity Prices with Nonlinear Models." Discussion Paper No. 268. Helsinki: Helsinki Center of Economic Research, University of Helsinki.

Allen, P.G. 1994. "Economic Forecast in Agriculture." *International Journal of Forecasting* 10 (1): 81–135.

Alonso, J.C., and Arcila, A.M. 2012. "Un modelo de predicciones diarias para contratos de futuros del azúcar." *Economía y Región* 6 (2): 33–51.

Apergis, N., Artikis, P. G., and Kyriazis, D. 2015. "Does Stock Market Liquidity Explain Real Economic Activity? New Evidence from Two Large European Stock Markets." *Journal of International Financial Markets, Institutions and Money* 38: 42–64.

Aradhyula, S., and Holt, M. 1988. "GARCH Time Series Models: An Application to Retail Livestock Prices." *Western Journal of Agricultural Economics* 13 (2): 365–374.

Barnett, W., Gallant, R., Hinich, M., et al. 1997. "A Single-Blind Controlled Competition among Tests for Nonlinearity and Chaos." *Journal of Econometrics* 82 (1): 157–192.

Belousova, J., and Dorfleitner, G. 2012. "On the Diversification Benefits of Commodities from the Perspective of Euro Investors." *Journal of Banking and Finance* 36 (9): 2455–2472.

Benavides Perales, G. 2009. "Price Volatility Forecasts for Agricultural Commodities: An Application of Historical Volatility Models, Option Implieds and Composite Approaches for Future Prices of Corn and Wheat." *Revista de Administración, Finanzas y Economía* 3 (2): 40–59.

Brace, M.C., Schmidt, J., and Hadlin, M. 1991. "Comparison of the Forecasting Accuracy of Neural Networks with Other Established Techniques." In: M.A. Sharkawi and R.J. Marks, editors. *Proceedings of the First International Forum on Applications of Neural Networks to Power Systems, Seattle, Washington July 23–26, 1991*. New York: IEEE.

Caire, P., Hatabian, G., and Muller, C. 1992. "Progress in Forecasting by Neural Networks." In: *Proceedings of The International Joint Conference on*

Neural Networks, vol. 2. IEEE.
Caraiani, P. 2012. "Nonlinear Dynamics in CEE Stock Markets." *Economic Letters* 114 (3): 329–331.
Caridad y Ocerín, J.M., and Caridad y López del Río, D. 2014. *Series temporales y predicción económica*. Spain: DF.
Chang, C., McAleer, M., and Tansuchat, R. 2012. "Modelling Long Memory Volatility in Agricultural Commodity Futures Returns." *Annals of Financial Economics* 07 (2): 1250010.
Coronado Ramírez, S.L., Ramírez Grajeda, M., and Celso Arellano, P.L. 2012. "Análisis sobre la no linealidad del rendimiento diario del precio internacional del café." *Revista Mexicana de Agronegocios* 31: 55–61.
Cromwell, J., and Labys, W. 1993. "Testing for Nonlinear Dynamics and Chaos in Agricultural Commodity Prices." Research Paper No. 9326. Research Institute, West Virginia University.
Deaton, A., and Laroque, G. 1992. "On the Behaviour of Commodity Prices." *The Review of Economic Studies* 59: 1–23.
Denton J. W. 1995. "How Good Are Neural Networks for Causal Forecasting?" *The Journal of Business Forecasting* 14 (2): 17–20.
Dicle, M.F., and Levendis, J. 2011. "Greek Market Efficiency and Its International Integration." *Journal of International Financial Markets, Institutions and Money* 21(2): 229–246.
Ferreira, L., Moura, G.L., Borenstein, D., et al. 2011. "Utilização de Redes Neurais Artificiais como Estratégia de Previsão de Preços no Contexto de Agronegócio." *Revista de Administração e Inovação* 8 (4): 6–26.
Freiman, J.P., and Pamplona, E.O. 2005. "Redes neurais artificiais na previsão do valor de *commodity* do agronegócio." Paper presented at the 5th *Encuentro Internacional de Finanzas*. Santiago, Chile.
Funahashi, K. 1989. "On the Approximate Realization of Continuous Mappings by Neural Networks." *Neural Networks*. 2 (3): 183–192.
Ganeta, R., Romeo, L. M., and Gil, A. 2006. "Forecasting of Electricity Prices with Neural Networks." *Energy Conversion and Management* 47 (13–14): 1770–1778.
Gorr, W.L., Nagin, D., and Szczypula, J. 1994. "Comparative Study of Artificial Neural Network and Statistical Models for Predicting Student Grade Point Averages." *International Journal of Forecasting* 10 (1): 17–34.
Gupta, J., and Sankalp, S. 2017. "The Impact of Global Financial Crisis on Market Efficiency: An Empirical Analysis of the Indian Stock Market." *International Journal of Economics and Finance* 9 (4): 225–252.
Gupta, N., and Singla, R. 2018. "Weak Form Efficiency of Gold Prices in the Indian Market." *International Journal of Engineering Technology Science and Research* 5 (3): 396–401.
Hamid, K., Suleman, M.T., Ali Shah, S.Z., et al. 2010. "Testing the Weak Form of Efficient Market Hypothesis: Empirical Evidence from Asia–Pacific Markets." *International Research Journal of Finance and Economics*. 58: 121–133.
Hamid, S.A. and Iqbal, Z. 2004. "Using Neural Networks for Forecasting Volatility of S&P 500 Index Futures Prices." *Journal of Business Research* 57 (10): 1116–1125.
Heravi, S., Osborn, D.R., and Birchenhall, C.R. 2004. "Linear Versus Neural Network Forecasting for European Industrial Production Series." *International Journal of Forecasting* 20 (3): 435–446.
Hippert, H.S., Pedreira, C.E., and Souza, R.C. 2001. "Neural Networks for Short-Term Load Forecasting: A Review and Evaluation." *IEEE Transactions on Power Systems* 16 (1): 44–45.
Hornik, K., Stinchcombe, M., and White, H. 1990. "Universal Approximation of an Unknown Mapping and Its Derivatives Using Multilayer Feedforward Networks." *Neural Networks*. 3 (5): 551–560.
Hou, Y., Liu, F., Gao, J., et al. 2017. "Characterizing Complexity Changes in Chinese Stock Markets by Permutation Entropy." *Entropy* 19 (10): 514.
Ikeora, J.E., Charles-Anyaogu, N.B., and Andabai, P.W. 2016. "The Weak Form Efficient Market Hypothesis in the Nigerian Stock Market: An Empirical Investigation." *European Journal of Business, Economics and Accountancy* 4 (6): 93–105.
Jain, A., and Kumar, A.M. 2007. "Hybrid Neural Network Models for Hydrologic Time Series Forecasting." *Applied Soft Computing* 7 (2): 585–592.
Kang, S. 1991. "An Investigation of the Use of Feedforward Neural Networks for Forecasting." PhD Thesis, Kent State University.
Karali, B., and Power, G. 2009. "What Explains High Commodity Price Volatility? Estimating a Unified Model of Common and Commodity-Specific, High- and Low-Frequency Factors." Paper presented at the 2009 Annual Meeting of the Agricultural and Applied Economics Association. Milwaukee, Wisconsin. July 29–29.
Khan, W., and Vieito, J.P. 2012. "Stock Exchange Mergers and Weak Form of Market Efficiency: The Case of Euronext Lisbon." *International Review of Economics and Finance* 22 (1): 173–189.
Kohzadi, N., Boyd, M. S., Kermanshahi, B., et al. 1996. "A Comparison of Artificial Neural Network and Time Series Models for Forecasting Commodity Prices." *Neurocomputing* 10 (2): 169–181.
Kristjanpoller, W., and Minutolo, M. C. 2015. "Gold Price Volatility: A Forecasting Approach Using the Artificial Neural Network–GARCH Model." *Expert Systems with Applications* 42 (20): 7245–7251.
Lean, H.H., Mishra, V., and Smyth, R. 2015. "The Relevance of Heteroskedasticity and Structural Breaks when Testing for a Random Walk with High-Frequency Financial Data: Evidence from ASEAN Stock Markets." In: G.N. Gregoriou, editor. *Handbook of High Frequency Trading*. London: Academic Press.
Liu, C. 2009. "Price Forecast for Gold Futures Based on GA-BP Neural Network." In: *2009 International Conference on Management and Service Science*. Piscataway, NJ: IEEE.
Lo, A.W., and MacKinlay, A.C. 1988. "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test." *The Review of Financial Studies* 1 (1): 41–66.
London Stock Exchange. 2018. http://www.londonstockexchange.com/.
Malliaris, A.G., and Malliaris, M. 2009. "Time Series and Neural Networks Comparison on Gold, Oil and the Euro." In: *2009 International Joint Conference on Neural Networks*. Piscataway, NJ: IEEE.
Mazviona, B.W., and Nyangara, D. 2013. "A Test of the Weak Form Efficiency of the Zimbabwe Stock Exchange after Currency Reform." *International Journal of Business, Economics and Law* 2 (2): 43–48.
Milas, C., and Otero, J. 2002. "Smooth Transition Vector Error Correction Models for Spot Prices of Coffee." *Applied Economics Letters* 9: 925–928.
Mynhardt, H., Makarenko, I., and Plastun, A. 2017. "Market Efficiency of Traditional Stock Market Indices and Social Responsible Indices: The Role of Sustainability Reporting." *Investment Management and Financial Innovations* 14 (2): 94–106.
Ntungo, C., and Boyd, M. 1998. "Commodity Futures Trading Performance Using Neural Network Models Versus ARIMA Models." *The Journal of Futures Markets* 18 (8): 965–983.
Ojeda Echeverri, C.A., and CastYEA Vélez, E. 2014. "Prueba de eficiencia débil en el mercado accionario colombiano." *Semestre Económico* 17 (35): 13–42.
Pinheiro, C.A.O., and Senna, V. 2017 "Forecasting of Prices through Neural Networks and Spectral Analysis: Evidence for the Future Market Commodities Sugar and Soy." *Custos e Agronegocio* 13 (4): 103–128.
Ribeiro, C.O., and Oliveira, S.M. 2011. "A Hybrid Commodity PriceĆForecasting Model Applied to the Sugar–Alcohol Sector." *The Australian Journal of Agricultural and Resource Economics* 55 (2): 180–198.
Ruiz-Gándara, A., and Caridad y Ocerin, J. M. 2014. "Computational Methods in the Identification of Forecasting Time Series Models." *International Journal of Scientific Management and Tourism* 1: 5–17.
Sallehuddin, R., Shamsuddin, S.M., Hashim, S.Z.M., et al. 2007. "Forecasting Time Series Data Using Hybrid Grey Relational Artificial Neural Network and Auto Regressive Integrated Moving Average Model." *Neural Network World* 17 (6): 573–605.
Samuelson, P.A. 1965. "Proof That Properly Anticipated Prices Fluctuate Randomly." *Industrial Management Review* 6 (2): 41–49.
Serin, P. 2017. "Is Random Walk Hypothesis a Reasonable Data Generating Process Assumption for Stock Prices?" Master's Thesis, Nord University.
Shaik, M., and Maheswaran, S. 2017. "Market Efficiency of ASEAN Stock Markets." *Asian Economic and Financial Review* 7 (2): 109–122.
Soon, L.P., and Abdul-Rahim, R. 2017. "Efficiency of Malaysian Stock Market: A Revisit Based on Analysts' Recommendations." *Geografia-Malaysian Journal of Society and Space* 12 (2): 1–14.
Srinivasulu, S., and Jain, A. 2006. "A Comparative Analysis of Training Methods for ANN Rainfall-Runoff Models." *Applied Soft Computing* 6 (3): 295–306.
Tang, S., Huang, S., Bowman, E., et al. 2017. "The Efficient Markets Hypothesis Does Not Hold When Securities Valuation Is Computationally Hard." Available at: https://ssrn.com/abstract=2925071.
Taskaya-Temizel, T., and Casey, M.C. 2005. "A Comparative Study of Autoregressive Neural Network Hybrids." *Neural Networks* 18 (5–6): 781–789.
Tejeda, H., and Goodwin, B. 2009. "Price Volatility, Nonlinearity and Asymmetric Adjustment in Corn, Soybean and Cattle Markets: Implication of Ethanol-Driven (Market) Shocks." Proceedings of the NCCC-134 Conference on Applied Commodity Price Analysis, Forecasting, and Market Risk Management. St. Louis, MO.
Tseng, C.H., Cheng, S.T., Wang, Y.H., et al. 2008. "Artificial Neural Network Model of the Hybrid EGARCH Volatility of the Taiwan Stock Index Option Prices." *Physica A: Statistical Mechanics and Its Applications* 387 (13): 3192–3200.
Tu, J.V. 1996. "Advantages and Disadvantages of Using Artificial Neural Networks versus Logistic Regression for Predicting Medical Outcomes." *Journal of Clinical Epidemiology* 49 (11): 1225–1231.
Velásquez Henao, J.D., and Aldana Dumar, M.A. 2007. "Modelado del precio del café colombiano en la bolsa de Nueva York usando redes neuronales artificiales." *Revista Facultad Nacional de Agronomía, Medellín* 60 (2): 4129–4144.
Villada, F., Cadavid, D.R., and Molina, J.D. 2008. "Pronóstico del precio de la energía eléctrica usando redes neuronales artificiales." *Revista de la Facultad Universitaria de Antioquia* 44: 111–118.
Villada, F., Muñoz, N., and García-Quintero, E. 2016. "Redes neuronales artificiales aplicadas a la predicción del precio del oro." *Información tecnológica* 27 (5): 143–150.
Wiles, P.S., and Enke, D. 2014. "Nonlinear Modeling Using Neural Networks for Trading the Soybean Complex." *Procedia Computer Science* 36: 234–239.
Yonenaga, W.H., and Figueiredo, R.S. 1999. "Previsão do preço da soja utilizando redes neurais." Paper presented at the 19th *Encontro Nacional de Engenharia de Produção* (ENEGEP). Rio de Janeiro.
Yu, S., and Ou, J. 2009. "Forecasting Model of Agricultural Products Prices in Wholesale Markets Based on Combined BP Neural Network-Time Series Model." In: *2009 International Conference on Information Management, Innovation Management and Industrial Engineering*. Los Alamitos, California: IEEE.
Zhang, G.P. 2004. *Neural Networks in Business Forecasting*. Hershey: Idea Group Publishing.
Zhang, G.P., Patuwo, B.E., and Hu, M.Y. 2001. "A Simulation Study of Artificial Neural Networks for Nonlinear Time-Series Forecasting." *Computers & Operation Research* 28 (4): 381–396.
Zou, H.F., Xia, G.P., Yang, F.T., et al. 2007. "An Investigation and Comparison of Artificial Neural Network and Time Series Models for Chinese Food Grain Price Forecasting." *Neurocomputing* 70 (16–18): 2913–2923.

CASE STUDY

Agriculture 4.0

Adding value to the primary sector

Twice the output. That is the seductive promise that artificial intelligence (AI) has made the agriculture sector by helping businesses to be managed more profitably. Agriculture is one of the main export sectors for Latin America, so many companies have developed tools and algorithms offering new services to primary producers. The Uruguayan company OKARATech has pioneered the creation of software platforms that combine big data techniques with AI. The system, which can be accessed from mobile devices, obtains, processes, and consolidates data from different sensors and information sources and makes this available to users. The company's CEO, Leonardo Cristalli, analyzes the current and future impacts of AI on agroindustry.

Is the agricultural sector going through a digital transformation?

Among the technologies that will be key to this sector are biotech, software and AI, big data, simulation, geostatistics, and robotics, including drones and autonomous vehicles. AI is particularly crucial because it can play a part in duplicating the sector's global output. Most agricultural exporters grow crops using production schemes and formulas that are based on average inputs, which are applied as a one-size-fits-all recipe throughout an entire region or property, with no differentiation between zones. This can lead to inappropriate use of inputs and resources. We have a very advanced capacity for understanding how biological systems are impacted by all kinds of variables—meteorological ones, geological ones, and others related to farming environments. By using technology, we can optimize this understanding even further. Today we have the computational power to provide high-speed responses (hardware) and tools to interpret the information (software and AI algorithms). This technology processes all the variables that are in play and interprets the environment to arrive at the most appropriate combination of inputs to apply to a given area, develop models that maximize production, simulate different scenarios, and make predictions.

What are the main benefits for growers?

Efficiency is key. From a digital perspective, the aim is to make more stringent use of scarce resources and even replace conventional energy sources with the use of solar or wind power. Another aim is to use autonomous vehicles, which are smaller, safer, and more precise. This will significantly change the dynamics of the business and how people go about it. By using these platforms, businesspeople can improve the decisions they make around inputs and financing schemes.

Where does OKARATech fit in?

To generate value, what is needed is an integrator to consolidate and standardize all the data from different sources, sensors, and proprietary systems, including multispectral images, soil samples, precision machinery data, and meteorological variables. First, our platform processes and standardizes that data and then offers it as a service that can be accessed from all kinds of devices and operating systems. The second phase is about using that data and interpreting it depending on the operating conditions within each user's business. We give producers a wide range of variables to base their decisions on.

Which recent developments have been most revolutionary?

There have been some dramatic developments in natural language processing, image recognition, and predictive capacity. In these fields, technologies and developments that large companies like Apple, Amazon, or Google have generated for end users are now being included in business-to-business (B2B) environments. Our focus is on adapting those technologies to agroindustry. These technologies are expected to mature in the next few years. In the future, for example, we hope that commercially available vision-based algorithms will be able to detect a plague or disease and can then warn the growers to spray only the affected area. The agriculture of the future will be digital, and these tools will be the common denominator. Indeed, without these tools, the system will not be viable. But the sectors still needs to experience significant cultural change and a paradigm shift.

> EFFICIENCY IS KEY. THE AIM IS TO MAKE BETTER USE OF SCARCE RESOURCES AND REPLACE CONVENTIONAL ENERGY.

Is there resistance to the use of new technologies in the agricultural sector?

Although there are significant differences between generations of farmers, many still make very limited use of cell phones and don't even use email. However, many others can no longer imagine their lives and their work without these tools. Training is another challenge: university agronomy programs are out of date: they should be teaching students to code and use geographic information systems (GIS) platforms. These contents currently come from the private sector, not academia.

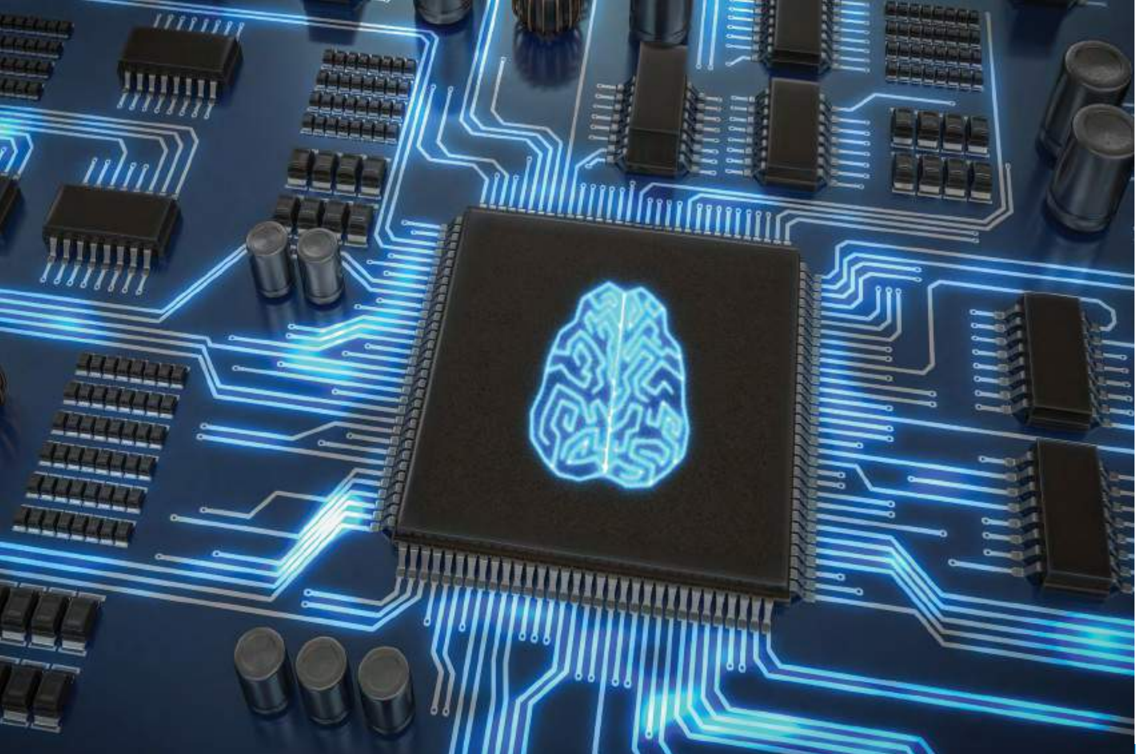

ALGORITHMIC TRADING

A Revolution in Capital Markets

Julián R. Siri and Juan Andrés Serur
University of CEMA

HOW CAN ARTIFICIAL INTELLIGENCE BRING DOWN THE COST OF FINANCIAL TRANSACTIONS AND PROVIDE COMPETITIVE ADVANTAGES TO MARKET AGENTS WHO USE NEW TOOLS? HOW CAN AI BE USED TO ANALYZE INVESTOR PROFILES AND DETECT SYSTEMIC RISKS THAT ARE AMPLIFIED AS STOCK MARKET TRANSACTIONS ARE AUTOMATED?

Decades before artificial intelligence (AI) became popular with mass audiences,[1] capital markets and stakeholders in these (investment banks, hedge funds, and brokerage firms, among others) began testing breakthrough technologies in this fledgling field that was fast becoming a reality. The innovative and highly competitive nature of the key players in the finance industry is what makes the need to survive and remain on the crest of the wave a deciding factor in the industry's receptiveness to these new technologies.

While most people are worrying about the potential impact of automation on employment and wages, capital markets have been making machines work to their benefit in areas as varied as algorithmic trading, quantitative analysis, mass data processing, and even robo-advisory services.[2]

Understanding AI as being not just a technology for bringing down costs but a tool that can generate value along the length and breadth of organizational structures has put early adopters in a privileged position within this new ecosystem.

As a science that could rewrite the rules of the game, AI has gone through many booms and busts. However, given recent technological advances, the refinement of methodologies, and the exponential growth in development-focused investment, everything seems to suggest that AI has now got a firm footing in the real world. Many specialists are calling this phenomenon the "Fourth Industrial Revolution" or the "Artificial Intelligence Revolution" and are predicting productivity increases on par with those of the first Industrial Revolution, over two centuries ago.

METHODOLOGIES, TRENDS, AND CURRENT AFFAIRS

For a long time, scientists working in different fields have tried to create models to enable machines to perceive their surroundings, understand problems, predict behavior, and come up with solutions for learning and improving quality of life. However, with the advent of AI, pioneers in the field set their sights on an even more ambitious goal, one that is not limited to understanding the world around us but which instead seeks to generate algorithms that can perceive, learn from, and understand their surroundings and solve problems that would at first appear to need to draw on human intelligence. From solving complicated puzzles, creating autonomous vehicles, and even being able to diagnose illness at an early stage, these technologies are pursuing the ambitious aim of not only thinking like humans but also making a qualitative leap in the quality of decision-making so as to contribute to the quality of everyday life.

Although AI has grown exponentially in recent years, it is much older than is generally believed. The term "artificial intelligence" was coined in 1955 by John McCarthy, a professor at Dartmouth College. He argued that AI is the science and engineering of creating intelligent computers. Different schools of thought have emerged within the field of AI throughout its history. There are two main trends: on the one hand, the school that supports the idea that AI is about machines and systems that are capable of adopting human behavior and, on the other, the school that defends rational machine behavior as being the definition of AI. The split mainly lies in whether computers are able to think and talk like humans, and have minds in a complete, literal sense (Bellman, 1978; Haugeland, 1985), versus the notion that computers think and act rationally (Charniak and McDermott, 1985; Schalkoff, 1990; Winston, 1992; Luger and Stubblefield, 1993).

The most widely accepted contemporary definition was put forward by Poole, Mackworth, and Goebel (1998) and clearly explains what AI is all about. The authors claim it to be the theory and development of computer systems that are able to carry out tasks that are normally restricted to the realm of human intelligence. As intelligent agents, human beings can perceive their environment, interpret conditions, and make decisions that maximize the probabilities of success at the target tasks. Machines attempt to replicate the cognitive abilities of the human mind, such as visual perception and voice recognition, in order to learn and solve ever more complex problems.

Throughout history, scientists and researchers in the fields of computing and statistics have developed techniques to achieve these goals, often by wringing out huge, dissimilar datasets. There seems to be no limits to this: any kind of data can be used, including data from different sources and of different qualities. These techniques can balance out computers' ability to undertake tasks such as recognizing images or processing natural languages by learning from experience.

Generally speaking, AI systems can be understood as an iterative process with the following features.

1. Acquiring information. Here the ability of AI focuses on recognizing and acquiring data in structured formats (such as economic data) and unstructured ones (images and sound).
2. Interpreting data. Systems are able to analyze data to reach conclusions or knowledge that is relevant to the challenge at hand.
3. Acting accordingly. AI can use this understanding of information to carry out a process, activity, or defined function.
4. Learning. Based on the feedback it receives from experiments carried out in the real world, AI is able to adapt and improve its effectiveness and efficiency over time. This particular feature is essential to distinguishing AI from routine automated processes.

However, many terms are used to describe this area of knowledge, so it makes sense to review these here and define and describe the scope of each. Although there is no clear and universally consistent definition for "big data," we use the one suggested by Ward and Barker (2013), who say that "big data is a term describing the storage and analysis of large and or complex data sets using a series of techniques," which include AI. The analysis of these sets is usually called "big data analytics," and the complexity of this is proportional to the quantity of unstructured or semi-structured data contained in the data sources.

THE LATIN AMERICAN INTEGRATED MARKET (MILA)

May 2011, after years of conversation and dispute over the agreements that would govern it, the Latin American Integrated Market (MILA) was established, and currently consists of the stock markets of Chile, Colombia, Mexico, and Peru.

An integrated market like this brings opportunities for investors to open up their portfolios to assets from other countries while also increasing available liquid assets and boosting cash flow capture for issuing companies, thus expanding the capital market into all MILA member countries. Integrated markets imply larger numbers of issuers, assets, and investors and thus help generate appropriate conditions for promoting the use of AI-centered technologies.[3]

With a view to the future of the investment management industry in general and focusing on MILA in particular, AI opens the door to potential benefits from new robo-advisory services, which entail the partial or total automation of fund management services, seeking to reduce entry barriers to capital markets for retail investors. The development of these sorts of applications has been enabled by both the adoption of direct market access (DMA) technologies and the standardization of regulations around investment vehicles and tax treatments.

These advances have enabled the creation of automated recommendation systems, which do everything from investor profiling to creating and rebalancing portfolios. Unsupervised learning techniques can be extremely useful in generating independent recommendations and guidance systems for investors.

When this type of system is being developed, the information that potential investors can provide is essential. One way of beginning the process is to analyze the attributes that are generally used when the investment policy plan is being created.[4] Large numbers of investor profile records, combined with machine learning techniques, can add value by improving the classifications of new investors that enter the system. Subsequently, algorithms can create investment portfolios, taking into account the universe of available assets and the appropriate risk profile, which were obtained during the first stage of the process. Algorithms can even be responsible for implementing and periodically rebalancing the portfolio. To complete the production cycle, algorithms can even manage the periodic reports sent to investors/users and other matters relating to subscriptions and fund redemptions.

We understand AI as a broad science, which so-called machine learning is a subcategory of. While AI is the theory and development of competing systems that are able to carry out tasks that would normally call on human intelligence, machine learning is a methodology that refers to the design of a series of actions to solve a problem (a process which is known as an algorithm), which is repeatedly optimized through experience that is collected during the process, under a certain degree of human supervision (or sometimes none at all). This was stated explicitly by Arthur Samuel (1959), one

US$ 7 BILLION

WILL BE INVESTED IN AI FOR FINANCE BY 2019

of the fathers of this branch of AI, who defined machine learning as a “field of study that gives computers the ability to learn without being explicitly programmed.”

Machine learning problems are generally made up of an error function, a loss function, and a target function. The aim is to use this learning process to minimize losses and achieve the specified goal.

There are two categories of machine learning, both of which are based on the level of human intervention needed to tag data and the complexity of the problem-solving techniques and the type of data used.

In the case of supervised machine learning, algorithms receive a set of training data containing tagged information (for example, transactions that have been identified as fraudulent). By doing so, part of the dataset that feeds the algorithm includes a direct, positive classification. The algorithm thus learns a general classification which it will use to predict the tags for the rest of the uncatalogued entries (to continue with the above example, this would entail determining which transactions are fraudulent and which are not). The learning process finishes when the algorithm achieves a reasonable level of precision in relation when carrying out its target function. This category includes classification and regression problems. In the first case, the target variable is discrete/categorical (a given transaction either is or is not fraudulent). In the second case, the target variable is continuous (transaction amount)

Unsupervised learning is another type, in which the information provided has no tags to classify it; in other words, the algorithm is not given the correct answer for each observation. This means that during the learning process, the relevant variable is not available. What the algorithm has to do is to detect patterns by identifying groups of observations with similar attributes and features. Although there are different techniques for exploring unsupervised learning problems, the most common is cluster analysis, in which the algorithm looks for common features in the data to create tag groups for each entry.

Halfway between supervised and unsupervised learning is reinforcement learning. The algorithm receives a set of unclassified data, picks a given action for each specific entry, and gets feedback (potentially from a human), which helps it to learn and improve. Generally speaking, reinforcement learning requires an exploration stage (most of the available data tends to be used at this point) and an exploitation stage, which seeks to draw on the learning acquired during the earlier stage.

The last category is deep learning. This type of AI is seeking to solve one of the greatest challenges to traditional machine learning models: extracting meaningful features. In the other models described above, the programmer needs to tell the computer what types of things it should use as data when making decisions. Deep learning algorithms are one of the few methods through which this challenge can be avoided. The algorithms can learn to focus on the right characteristics by themselves and require little orientation from programmers. These algorithms work using layers that were inspired by the structure of the human brain and thus are called neural networks. The neurons in the network have weighted inputs that are filtered by a specific function. The value that the function arrives at is transmitted as an output to another neuron. A neural network is created when we connect the neurons to one another and to the input and output data that the algorithm is trying to respond to. Just as the neurons in the human brain are organized into layers, so are the ones in these artificial networks. The neurons in the lower layers receive input signals while the ones in the higher layers receive signals from the layers below them. The last layer of neurons connects to the answers. The underlying objective is to process and classify complex datasets whose nonlinear behavior prevents the use of simpler algorithms.

It is important to understand the context for this boom in the implementation of AI methodologies. The availability of huge quantities of data is what has enabled scientists to develop, train, and apply these algorithms. In recent years, all kinds of data sources have emerged: from comments and publications on social media (including audio and video content) to satellite imagery or online shopping receipts. This all eventually translates into billions of gigabytes of new information that is being produced day in, day out. If we add to this the fact that data storage costs have been plummeting and data analysis capacity is growing exponentially (on both the software and hardware fronts), it all adds up to an ecosystem that is predisposed to the degree of innovation that we are currently witnessing.

With regard to the level of AI development within the world of investment, the quantities of capital pledged to startups within the sector is an undeniable sign of the prevailing optimism, especially in finance and banking. According to CB Insights[3], since 2012, more than US$15 billion has been invested

FIGURE 1
AI-RELATED M&A ACTIVITY

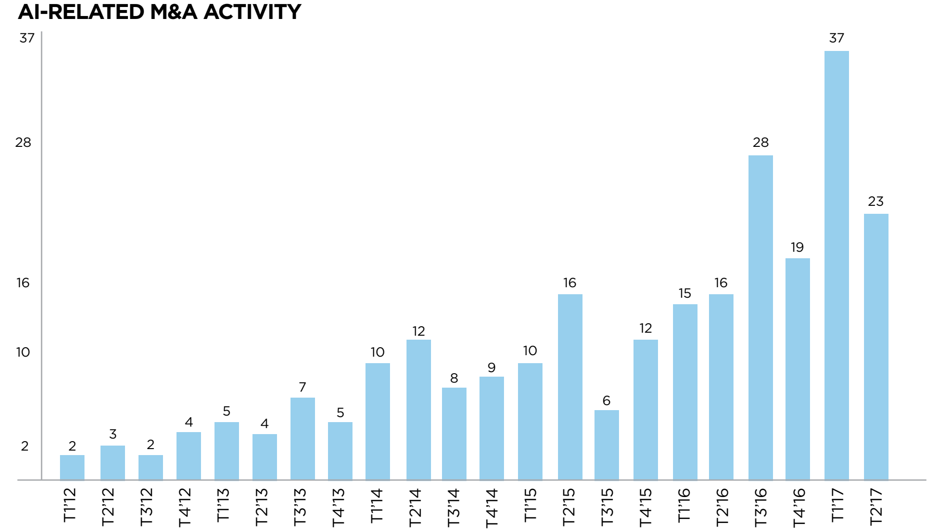

Source: CB Insights (2018).

in funding over 2,300 AI projects. The money pouring into development has been growing at a year-on-year rate of over 70%. In the finance industry, it is estimated that investment in AI will grow from US$2 billion in 2016 to more than US$7 billion by the end of the decade.

The number of academic publications that focus on AI has increased nine-fold since 1996, a growth rate that is disproportionate to that of publications in other branches of science. For example, the number of computer science publications has increased six-fold in the time period mentioned above while the number of publications in other areas of science has only doubled.

Another bellwether for the levels of enthusiasm associated with AI is the fact that mergers and acquisitions among companies associated with these technologies have been consistently on the rise. For example, Alphabet Inc. alone acquired more than 50 such startups between 2015 and 2016.

APPLYING AI TO CAPITAL MARKETS

AI has been adopted quickly in capital markets and is used for a wide range of practical applications. The reasons given for the speed of its adoption include factors on both the supply and demand sides.

Those on the supply side include technological progress, by which we mean ever more powerful processors, lower hardware costs, and the appearance of cloud computing, a disruption that enabled the mass scalability needed to make AI a reality. Not only have storage and data analysis costs plummeted to a tenth of their value a decade ago, the volumes of information generated have increased seven-fold during the same period (Reinsel, Gantz, and Rydning, 2017; Klein, 2017). Two factors that have contributed to the rise in the use of these technologies include the emergence of larger quantities of financial (and nonfinancial) data that can be applied to the problems discussed here and the fact that this data can now be accessed by growing numbers of people.

On the demand side, financial institutions, given the nature of their business, are natural seekers of IA-related technologies. Profitability is a key issue, in that these technologies bring opportunities to reduce costs and improve risk management. Other notable motivating factors include the fierce competition between banks. Likewise, the need to develop services associated with these new technologies can affect their reputations. Finally, demand has increased due to regulatory requirements, in that greater regulation has given rise to the need for efficient compliance. This has prompted banks to automate all reporting processes to a greater or lesser extent.

As a report from the Financial Stability Board (2017) stresses, use cases of AI in the sector include a range of areas and specific focuses. These range from applications that focus on the end user of banking services (credit scoring, valuation and sale of insurance policies, and digital assistants or bots to respond to customer requests) to others that entail decision-making within capital markets. We will be analyzing the latter in greater detail from an operational perspective and looking at their role in generating trading ideas and structuring investment portfolios.

Banks and investment funds use AI tools from an operational point of view for tasks such as optimizing the use of

INVESTOR CLASSIFICATION

One possible use of machine learning algorithms is to improve the classification of investors by their characteristics: pigeonholing them based on their risk-return profile is key to implementing investment strategies for their assets. Specifically, the K Nearest Neighbors (KNN) algorithm is an efficient tool for segmenting investors according to their level of risk tolerance. This is extremely useful as it can be used by robo-advisory services providing automated handling of investment portfolios.

The first step consists of defining the variable to be predicted (target variable) and the explanatory variables (features).

Target variable: investor risk tolerance, with three alternatives: high, medium, low.

Explanatory variables (features): age, time horizon, gender, and income.

The aim is to train the algorithm using a tagged dataset and then test it with an untagged dataset.[7] The variables are standardized to achieve a certain homogeneity in their values:

$$X_i^{Norm} = \frac{X_i - X^-}{X^+ - X^-}$$

where X_i^{Norm} is the standardized variable i (which takes a value of between 0 and 1) , X_i is the variable to be standardized, X^- is the minimum value of the set of variables, and X^+ is the maximum value. Nominal values are transformed using dummy variables.

Using the variables described above, a score Z_i is calculated for each investor, on the basis of which their risk tolerance is defined:

$$\text{Tolerance } (Z_i) = \begin{cases} \text{High}, & Z_i > 0.66 \\ \text{Medium}, & 0.66 > Z_i \geq 0.33 \\ \text{Low}, & Z_i < 0.33 \end{cases}$$

Then, of a total 500 pieces of data from different investors, 350 are picked to train the algorithm and 150 are used to test it. The results are shown in the following table.

	HIGH	LOW	MEDIUM
OBSERVATIONS	43	55	52
ALGORITHM	39	57	54

It shows almost 95% efficiency (see figure 2).

FIGURE 2
KNN GROUP CLASSIFYING ALGORITHM
PREDICTION BASED ON 2 VARIABLES (FROM A TOTAL OF 4)

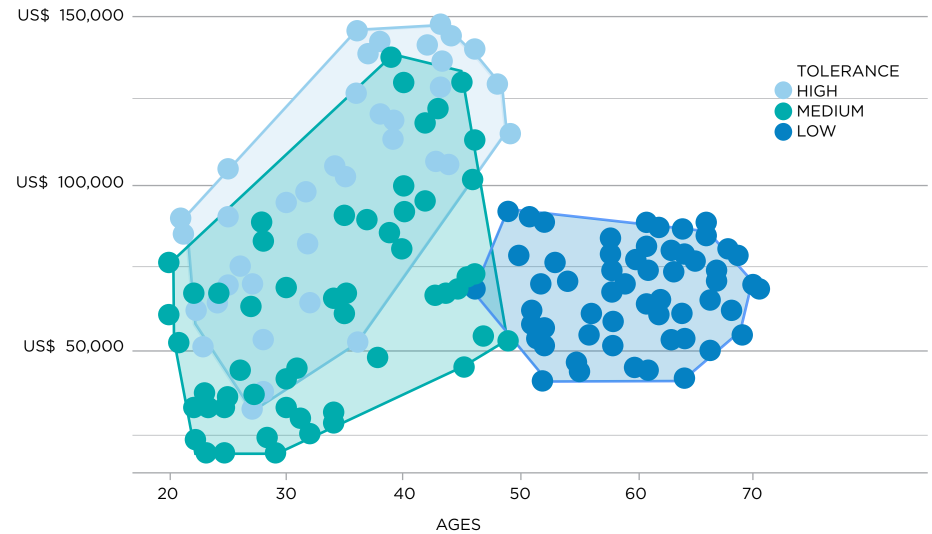

Source: Compiled by the authors.

capital, managing risk models, and analyzing and streamlining market impact (when putting together or taking apart large market positions with uncertain levels of liquidity). With regard to capital allocation, AI algorithms are used in conjunction with the regulations on required minimum margins to develop and implement strategies that reduce these margins and free up capital so as to improve profitability and the allocation of scarce resources. In risk management, new technologies focus on issues such as model validation, in which unsupervised learning algorithms are applied to detect anomalous projections generated by stress-testing models. By so doing, they provide support for the personnel responsible for validation by determining if risk management models are operating within an acceptable range of tolerance. Finally, in connection with market impact, algorithms capable of capturing the nonlinear relationships within order flow trading are used to find similarities among assets (within the same class) and order patterns to generate better estimations of price movements within markets. These are no minor issues when it is estimated that up to two-thirds of gross profit for quantitative funds (which execute systematic buy and sell orders) can be lost through costs associated with market impact (Day, 2017).

AI and machine learning techniques have gained ground in the field of trading and portfolio management, where they are used not just for signal generation but also to implement strategies. In trading, large brokerage firms generate vast flows of information on their operations. To analyze trading patterns and anticipate the client's next order, analysis and recommendation tasks need to be delegated to machine learning algorithms not just because doing so is much faster, but also because they have greater capacity for processing and interpreting the huge quantities of information that are constantly being generated. Based on risk models and algorithms, organized markets can also determine changes in user profiles, which would prompt an appropriate response or intervention. Within portfolio management, AI methodologies and tools allow new signals on price movements to be identified, due to their inherent capacity to use large databases of dissimilar information sources. Not only do they capture the analytical principles of what is known as systematic or rules-based investment, they also exploit patterns in price movements and even generate linear (and nonlinear) relationships as outputs of deep learning algorithms. At the end of the day, the aim of all these approaches is to generate greater returns on investment for different time horizons, with the lowest possible degree of correlation between strategies.

BENEFITS AND CHALLENGES

Like any experimental field in the full upward swing of its maturity cycle, AI brings benefits but also poses different sorts of challenges that need to be taken into account if we are to take advantage of it appropriately.

The benefits of adopting these new technologies include the fact that more efficient data processing should contribute to generating more efficient financial systems while improving perceptions of the stakeholders involved in this, reducing information asymmetries, and fine-tuning the valuation of the instruments used. AI applications need to be understood as tools that will gradually contribute to compliance with regulatory frameworks, thus improving the effectiveness of supervision. In a world in which concentration and integration within this industry are looked on with a certain amount of distrust (the "too big to fail" effect), having more, better-informed regulators is an advantage.

However, it is essential for us to contemplate potential challenges or problems that may be associated with the indiscriminate adoption of these technologies. We understand that there are risks at both the micro and macro levels. Moving from the general to the specific, as yet there are no international standards that establish a framework of reference in areas such as algorithmic trading. Many argue that machines and their opaque models[6] may potentially amplify systemic risks. Likewise, the lack of training in algorithm modeling and implementation and in interpreting the results from complex processes (which may exceed the speed and capacity of the human brain) are warning signs that we cannot blindly place our trust in such techniques. Furthermore, the scalability of these technologies and the so-called network effect (in the sense that they generate ever more sophisticated interconnections) may give rise to dependency on external suppliers. This could lead to the emergence of new systemically important players that lie outside the regulatory perimeter but that have the same impact on markets as traditional players.

APPROPRIATE RISK SUPERVISION

As is the case with almost any new financial product or service, AI must pass a series of demanding tests in relation to risk management and supervision. Given that we are discussing a topic as sensitive as money and an industry whose actions significantly affect the rest of the economy, it will be of paramount importance to evaluate the uses of AI from a broad perspective, one that adheres to the relevant protocols on data privacy, cybersecurity, and other risks.

It will also be key to test and train tools using unbiased data, emphasizing feedback mechanisms to ensure that applications function correctly. The future looks bright but balancing the unchecked greed that is sometimes present in capital markets is vital if the full potential of AI is to be unleashed.

NOTES

[1]For example, a Google Trends (https://trends.google.com) search for keywords such as "artificial intelligence" and "machine learning" will provide a proxy for Internet users' interest in AI and related issues.
[2]Robo-advisors are an automated investment management service that delegates the composition of investment portfolios, the designation of risk profiles and even the carrying out of orders to sophisticated AI algorithms.
[3]Online database reporting on angel investors and venture capital firms (https://www.cbinsights.com).
[4]As was mentioned above, one feature that is common to all AI algorithms to a greater or lesser extent is the need for large amounts of data to provide optimum results.
[5]Within this plan, the potential client provides a series of socio-economic, target, and risk-tolerance data.
[6]So-called black box algorithms (the workings of which are opaque or unknown) are seen as being undesirable within machine learning for the financial sector. When asked to detect patterns rather than causalities, these algorithms often act irrationally.
[7]In the training set, the data for both the explanatory variables (features) and the variable to be explained (target) is tagged. Furthermore, the untagged set (test set) includes explanatory variables but not target variables. The aim of doing so is to prove how efficient the algorithm is at classifying investors by their preferences based on the set of variables included.

REFERENCES

Artificial Intelligence Index. 2017. *2017 Annual Report*. AI Index.
Bellman, R. 1978. *Artificial Intelligence: Can Computers Think?* San Francisco: Boyd & Fraser Publishing Company.
Citibank. 2018a. *Banks and AI: The Compliance and RegTech Opportunity*. Global Banks Insights.
—. 2018b. *AI and Finance*. Global Banks Insights.
CB Insights. 2018. "The Race for AI: +Google, Intel, Apple in a Rush to Grab Artificial Intelligence Startups." *CB Insights*. February 27.
Charniak, E., and McDermott, D. 1985. *Introduction to Artificial Intelligence*. Boston: Addison-Wesley Publishing Company.
—Day, S. 2017. *Quants Turn to Machine Learning to Model Market Impact*. Risk Magazine. April 5.
Financial Stability Board. 2017. *Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications*. Basel: FSB.
Finextra and Intel. 2017. *The Next Big Wave: How Financial Institutions Can Stay Ahead of the AI Revolution*. London: Finextra.
Haugeland, J. 1985. *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press.
Institute of International Finance. 2017. *Deploying Regtech Against Financial Crime*. IIF.
J.P. Morgan. 2017. *Big Data and AI Strategies: Machine Learning and Alternative Data Approach to Investing*. Global Quantitative & Derivatives Strategy
Klein, A. 2017. "Hard Drive Cost Per Gigabyte." *Blackblaze blog*. July 11.
Luger, G.F, and Stubblefield, W.A. 1993. *Artificial Intelligence: Structures and Strategies for Complex Problem Solving*. Redwood City: Benjamin-Cummings Publishing Company.
Poole, D., Mackworth, A., and Goebel, R. 1998. *Computational Intelligence: A Logical Approach*. New York: Oxford University Press.
Reinsel, D., Gantz, J., and Rydning, J. 2017. "Data Age 2025: The Evolution of Data to Life-Critical." IDC White Paper. Framingham: IDC Headquarters.
Rich, E., and Knight, K. 1991. *Artificial Intelligence*. New York: McGraw-Hill.
Russell, S., and Norving, P. 2009. *Artificial Intelligence: A Modern Approach*. Third Edition. Harlow: Pearson.
Schalkoff, R.J. 1990. *Artificial Intelligence: An Engineering Approach*. Singapore: McGraw-Hill.
—. 1992. *Pattern Recognition: Statistical, Syntactic and Neural Approaches*. New York: John Wiley and Sons.
Samuel, A.L. 1959. "Some Studies in Machine Learning Using the Game of Checkers." *IBM Journal of Research and Development* 3 (3): 210–229.
Skilton, M., and Hovsepian, F. 2017. *The 4th Industrial Revolution: Responding to the Impact of Artificial Intelligence on Business*. Cham: Palgrave Macmillan International.
Ward, S.J., and Barker, A. 2013. "Undefined by Data: A Survey of Big Data Definitions." arXiv:1309.5821.
Winston, P.H. 1992. *Artificial Intelligence*. Reading: Addison-Wesley Publishing Company.

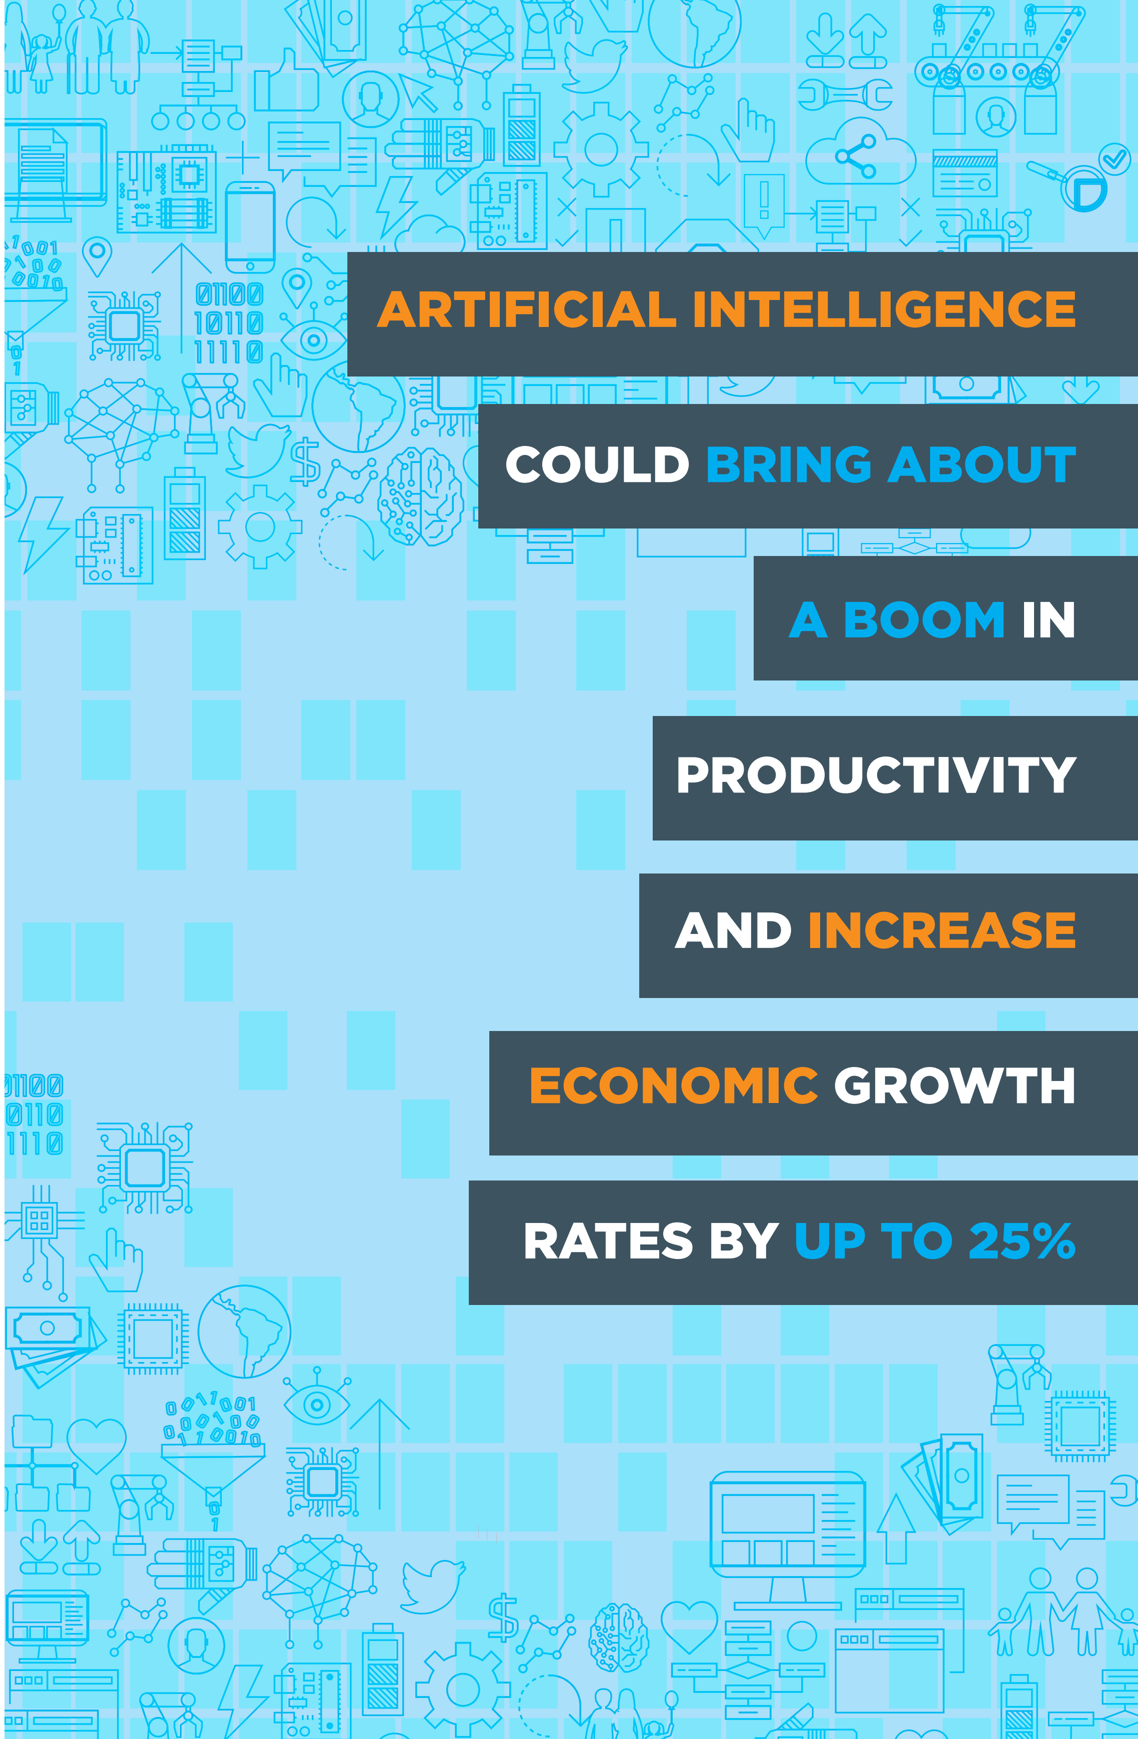

Neural Network Governance

New technologies pose unprecedented challenges for global governance, including ethical risks that require basic agreements to be reached. They also bring opportunities for improving public-sector efficiency and providing better services.

A Challenge to Global Governance

Nicolas Miailhe and Yolanda Lannquist
The Future Society

BETWEEN LATE 2017 AND EARLY 2018, THE FUTURE SOCIETY LAUNCHED A GLOBAL DEBATE ON ARTIFICIAL INTELLIGENCE (AI) THAT INCLUDED 2000 PARTICIPANTS AND 20 EVENTS AROUND THE WORLD. THIS ARTICLE HIGHLIGHTS THE MAIN POINTS THAT EMERGED FROM THIS DISCUSSION, LOOKS AT CURRENT AND FUTURE RISKS, AND PROVIDES A ROAD MAP FOR ARTICULATING A DECISION SYSTEM THAT WILL ENSURE THE IMPACT OF AI IS POSITIVE.

Anchored in the wider digital revolution, artificial intelligence (AI) is poised to transform the economy, society, and political systems we know today. Because of the network and scale effects around data, cloud supercomputing, and machine learning algorithms, the scale at which these dynamics are playing out is increasingly continental and global. The impact of the AI revolution combines very substantive opportunities and serious societal risks. The potential beneficial outcomes include a wave of productivity gains, hyper-tailored education, new drug discovery, safer roads, and efficient energy usage. However, AI also invokes the specter of widening inequality and mass unemployment, the threat of cyberattacks and lethal autonomous weapons, and the loss of human privacy, dignity, fairness, and agency. The prospect of shaping the development of AI development so as to capture the upsides while minimizing the downsides will primarily depend on the type of policy mixes countries deploy and their ability to collectively shape global governance processes.

Before analyzing the drivers of global governance processes for the age of AI, we need first to better understand the contours of the notion of "artificial intelligence" in the context of its evolving development, interactions, and impact on society. This definition then serves as a lens through which to analyze the complex socio-economic system dynamics involved in the global rise of AI. These dynamics necessitate novel frameworks for the global governance of AI.

This paper leverages insights from The Future Society's (2018) global civic debate on the governance of AI, which took place from September 2017 through March 2018 and gathered insights from a global community of over 2,000 participants and 600 active contributors in five different languages. Based on a proven open-innovation methodology, this pioneering experiment combined an online collective intelligence platform curated by a team of experts with a diverse series of over 20 online and offline events organized around the world. This curated conversation involved AI experts, practitioners, policymakers, academics, and citizens.

AGREEING ON A DEFINITION

A major challenge in governance and policy for AI remains the lack of consensus around how to define it. This undermines the potential for measuring its dynamics and impact and reaching agreement on this. Though AI is firmly embedded in computer science and has been an integral part of its rise since the 1940s, today the term is used refers to a wide range of technologies and methods.

Professor Stuart Russell, co-author of the seminal textbook Artificial Intelligence: A Modern Approach (2016), defines AI as the “study of methods for making computers behave intelligently.” This includes taking actions likely to achieve a specific end, or, in technical terms, to maximize expected utility. Consequently, Russell suggests that AI includes tasks such as learning, reasoning, planning, perception, understanding language, and robotics. AI is therefore an umbrella term that refers to an array of technologies that rely on algorithms at their core to “think” or “act” like humans. AI technologies include machine learning, computer vision, smart robotics, robotic process automation, biometrics, swarm intelligence, virtual agents, natural language generation, and semantic technology, among others.

Machine learning is the subset of AI driving recent developments in the field. This term, coined by Arthur Samuel (1959), refers to the “field of study within AI that gives systems the ability to learn from past examples to act in new and uncertain scenarios, without being explicitly programmed.” Machine learning lies behind AI innovations in fields as diverse as autonomous vehicles, personal assistance robots, chatbots, language translation, recommending Netflix films, and winning a game of Go. It encompasses several techniques, including neural networks and deep learning, reinforcement learning, regression analysis, clustering, decision trees, and more.

At their core, machine learning algorithms rely on statistics and mathematics to predict outcomes for new scenarios based on large training datasets. Written in code and powered by increasingly powerful computing systems, AI algorithms have developed alongside advances in computing infrastructure. While some machine learning techniques, such as neural networks, are loosely inspired by the complex networks of neurons that power the human brain, the convergence between computer science and neuroscience remains limited so far. Nevertheless, insights from the convergence between computer science and brain science through biotechnologies and bioinformatics may deepen in the coming decades as we develop greater understanding of the brain and its incredibly complex biochemical processes.

For example, the convergence of computer science and brain science may be one of several avenues with the potential to contribute to the development of artificial general intelligence (AGI). At present, AI systems have capabilities in specific discreet tasks, such as driving a vehicle or playing a game, which is referred to as artificial narrow intelligence (ANI). However, several scientists and technology entrepreneurs are currently working to develop AGI, which, for the majority of experts, refers to the ability to perform a full range of intellectual tasks comparable to those that the human brain can carry out. Experts widely disagree about the time horizon for AGI development, but they agree that it has the potential to trigger a large, paradigm-shifting impact on humanity.

FIGURE 1
AI AT THE INTERSECTION OF THREE TECHNOSCIENTIFIC MEGATRENDS

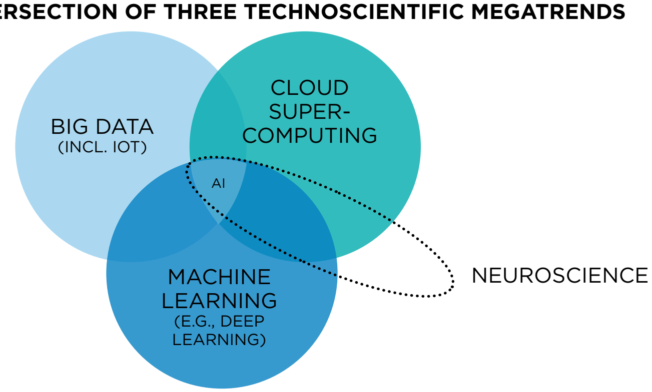

Source: Nicolas Miailhe, Yolanda Lannquist, The Future Society.

MEGATRENDS IN SOCIETY AND TECHNOLOGY

Machine learning algorithms learn to predict and act in new scenarios from thousands or even millions of data points, which include labeled images, past consumer purchases, or miles driven. Machine learning algorithms use this training data as input to learn from and predict outcomes for new data or new scenarios. This training process requires high levels of computing power. The result is the development of algorithms that perform very accurately when they are faced with new data or scenarios.

The recent renaissance in AI thus lies at the intersection of three technological megatrends: big data, machine learning, cloud supercomputing (Miailhe, 2018). Although AI has been a field of study for over half a century, the dizzying growth of computing power and the availability of large flows and stocks of digital data have boosted the development of machine learning (figure 1).

These three components do not lie within the vacuum of technology. They are defined, created, and implemented by humans. Engineers build large datasets; design, test and parameterize algorithms; interpret outputs; and determine how these are implemented in society. Humans are present at the design, input, operational, output, and implementation phases, and are deeply embedded in AI. We can therefore call AI a sociotechnical phenomenon. Equipped with ever smarter phones and devices, billions of people utilize, inform, and are affected by AI globally every day. Our digital lives fuel the development of AI as we continuously provide high-resolution digital data in the form of social media activity, transactions, and behaviors. AI is nested in the digital revolution, including the rise of social media, digital platforms and the digital economy, the Internet of Things (IoT), and cloud technology, all of which are sociotechnical systems that depend on humans, technology, and the interaction between the two.

Unlike than the 1950s dream, often depicted in films, of machine intelligence replicating common sense, consciousness, or emotions, AI complements but does not precisely replicate human intelligence as we know it. AI systems take into account more information and make decisions more rapidly and often more accurately than humans, without the interference of emotions, while developing parallel and often greater-than-human decision-making processes.

To establish a working definition, we can therefore define AI as being big-data–driven, machine-learning-algorithm–centric, sociotechnical systems powered by supercomputing. This definition serves as a relevant lens through which to assess the dynamics involved in the rise of AI and decide how to govern this global phenomenon.

THE RISE OF COMPLEX DYNAMICS

Competition among global firms and nation-states in the global race to develop AI may accelerate innovation at the expense of ethical and safety standards. Differences in regulatory and consumer protection regimes lead to potentially destructive imbalances between countries with higher risk appetite for growth, development, and innovation, and those aiming to protect citizens from potential abuses. Processes of global coordination and governance are thus needed to balance these complex dynamics and to raise the bar and avoid a race to the bottom in terms of social impact, safety, and ethics.

Governing the rise of AI revolves around the complexity of striking the right balance between activating and supporting beneficial innovations, on the one hand, and mitigating downside risks and minimizing adverse effects, on the other. Beneficial AI-related innovation includes better medical diagnostics, personalized education, and efficient natural resource allocation. AI can be implemented to make public services more efficient and accessible, provide safer transportation, achieve accurate medical diagnostics earlier, and democratize access to legal services for all sectors of society. There is critical value in the inclusive communities AI could create through providing access to such services. Significant gains in productivity and economic growth from new product offerings and improved supply chains promise to counteract the burden that aging populations in developed countries represent for the global economy. For example, McKinsey & Company estimates that implementing artificial neural networks in a range of business functions across 19 industries will derive in US$3.5 trillion to U$5.8 trillion per year in economic value.

However, without balanced ethical and safety standards, AI poses major societal risks. Cybersecurity, data integrity, and lethal autonomous weapons threaten our basic safety and security. In healthcare, there may be a trade-off in data privacy and security to achieve more accurate diagnostics and personalized treatment. In transportation, autonomous cars can reduce fatalities and carbon emissions while displacing millions of jobs. The collection of representative, accurate, and up-to-date data is often compromised for the sake of collecting greater volumes of data, leading to algorithms that are biased or unfit for some populations. This is a major concern for healthcare, justice systems, access to finance, and other sectors impacting lives. Ethics, liability, and value alignment for autonomous systems, or ensuring that actions in new scenarios align with human values and objectives, are unresolved challenges. This is a critical risk as AI capabilities continue to grow in the direction of AGI, as a system with a full range of human intellectual capabilities is significantly more challenging to secure and control, so the stakes are much higher.

RAPID TECHNOLOGICAL ADVANCES

The rapid pace of technological advances impedes timely, relevant regulations and policymaking. Developments in AI technology are converging with other emerging technologies and the digital economy to create a rapidly changing landscape where consumer uptake is fast. Lawmakers and governing bodies are often unable to stay ahead of these technological trends, let alone anticipate new ones. Policymakers and technology developers cannot predict how new technologies will impact society. Knowledge gaps in government and communication gaps between government and technology sectors further preclude relevant policymaking.

Moreover, rapid technological advances are making planning and preparing the public for technological transformation more difficult. The vast majority of citizens are ill-informed about AI technologies and unaware of their potential risks and impacts. Consumers are unaware of threats to data privacy and security, fake news, and disinformation. Long time horizons are needed to implement policies to prepare the workforce for automation, including education and skills training or social welfare policies, and affect outcomes.

In a world that is interconnected in real time, the consequences and imbalances caused by AI will create new challenges. The Industrial Revolution and the mechanization of agriculture in the 18th and 19th centuries, the advent of automobiles, computers, and the information age in the 20th century, and the rise of social media and the digital economy in the 21st century all involved ever larger populations, faster speeds, and greater impacts. Today, technological innovation takes place much faster, as does the incorporation of these changes into our daily lives. New technologies with consumer-friendly interfaces are accessible to citizens, populations are larger, and markets are more interconnected. Given the speed and scope of these impacts, governing the rise of AI involves new challenges and higher stakes.

A FACTORY OF MONOPOLIES

Without appropriate governance or policies to ensure that markets are competitive, suboptimal monopolistic or oligopolistic market dynamics may arise from AI innovation. Access to and, increasingly, ownership of data, computing power, and high-skilled talent to operationalize machine learning algorithms are the factors that determine market leaders. As Goldfarb and Trefler (2018) demonstrate, AI development enjoys economies of scale and scope. Large digital firms gather more consumer data, hire stronger talent, and have the resources to build vast and dedicated hardware and cloud supercomputing capabilities. Machine learning algorithms addressing some applications are transferable to others through a procedure called transfer learning. Therefore, the current global market domination by the US companies Google, Amazon, Facebook, Apple, and Microsoft (GAFAM) and the Chinese companies Baidu, Alibaba, Tencent, and Xiaomi (BATX) are, to a certain extent, a natural consequence of the components required for AI development. At present, the top machine learning talent, largest datasets, and the greatest stores of computing power are concentrated in the innovation hubs centered around these firms.

Oligopolistic and monopolistic markets disadvantage society. Less constrained by consumer preferences than in competitive markets, market-dominating players have fewer incentives for offering higher quality, safe, or ethical products or services at lower prices. Historically, competitive markets with multiple players have proven key for numerous beneficial innovations, which have resulted in higher quality of life, better public services, solutions for environmental sustainability, and other gains. However, the need to innovate as quickly as possible and for as low cost as possible, as is the case in today's globally competitive AI market, disincentivizes safe and ethical innovations. Meanwhile, the development of the first AGI may lead to a "singleton" (a single globally dominant state or commercial actor), a group of actors, or a superintelligent machine. This singleton may or may not have citizens' preferences at the core of its objectives.

A RACE TO THE BOTTOM?

A global race to develop new AI technology and AGI as quickly as possible is currently being run between firms and states. Although competition accelerates innovation, the lack of market incentives or global coordination may drive a race to the bottom in standards of ethics, safety, and upholding human values such as privacy, dignity, fairness, transparency, and more.

It seems that the rise of AI within the digital industrial landscape strengthens winner-takes-all market trends, whereby the companies offering the most advanced technologies quickly dominate the market because of economies of scale and network effects. As mentioned above, this is exemplified by the current global market domination of AI market leaders GAFAM, in the US, and BATX, in China. In this context, firms have incentives to grow rapidly, innovate, and put products onto the market. In the rush to get products to market, companies bypass necessary safety and testing procedures. Safety and ethical precautions may be reduced or skipped altogether, including sufficient testing before deployment, the use of large and representative datasets, precautions for data security against cyberthreats, and building value alignment and control in autonomous systems. According to AI safety expert Professor Roman Yampolskiy, "We have a pattern of preferring performance over safety, and that's what markets usually prefer."[1]

Meanwhile, national actors are increasingly aware of the strategic economic, political, and military issues at stake. Following the United States' lead in 2016, China, France, and European Union have recently pledged billions of dollars in investment and outlined concrete strategies toward the goal of becoming global AI leaders (Cerulus, 2018; European Commission, 2018). National economic interests are at stake, which have the potential to capture financial and economic gains from new product offerings and markets. AI is also impacting electoral politics: in the last decade, machine learning techniques have been used to target and personalize the political messaging disseminated through social media by official political campaigns and also by unofficial bots, fake news, and disinformation.

The strategic potential of AI at the intersection of defense, security, and computing could lead to a global AI arms race. The militarization of AI—including the development of lethal autonomous weapons, which can select and engage human targets without human control—could lead to game-changing shifts in the military power balance and regional or global instability.

The significant first-mover advantage in the militarization of AI disincentivizes actors from making time for precautions or abiding by regulations and standards. Beyond states, bad actors such as criminals, tyrants, and terrorists could hack such weapons, initiate cyberattacks using automated AI systems, or develop their own weaponized drones and vehicles to target the public. Policymakers in the United States and Europe largely agree on the need to ban lethal autonomous weapons, but the lack of global coordination and the limited capacity for monitoring and enforcing decentralized AI development pose major challenges.

Mechanisms to govern the safe development of AGI involve similar challenges. Researchers at Oxford University's Future of Humanity Institute have used game theory to model the global race for AGI development. They have found that an increase in the number of actors developing AGI and enmity between them both increases the danger of an AI disaster. "Under the assumption that the first AI will be very powerful and transformative, each team is incentivized to finish first—by skimping on safety precautions if need be" (Armstrong, Bostrom, and Shulman, 2013).

COMPETITIVE DISADVANTAGES

In the race to develop AI, firms and countries upholding safety and ethi-

FIGURE 2
THE GLOBAL AI RACE

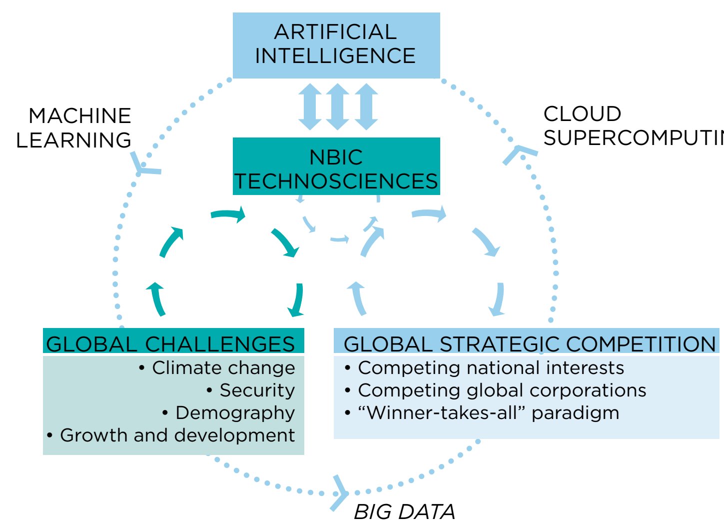

Source: Nicolas Miailhe, The Future Society cal procedures may fall behind. The regulation of development, such as data privacy or AI safety measures, can slow the rate and increase the costs of AI innovation. Limiting access to digital consumer data deprives machine learning algorithms of the valuable training data needed to improve performance. Compared to lax regulations for training and testing autonomous systems, pre-emptive certifications and standards for the safe implementation of AI and for fair and representative data may be detrimental to innovation.

Risk appetite in the trade-off between regulation and innovation varies across countries, as some government and industry leaders are readier to sacrifice citizens' safety than others. Cultural differences are another factor: 93% of Chinese customers are willing to share location data with their car manufacturer, compared to 65% of Germans and 72% of Americans (McKinsey & Company, 2016). For example, at present, China has an ambitious strategy to lead the world in AI development by 2030. To support this goal, it ensures that AI companies have access to vast stores of citizens' digital data (Ding, 2018: 25). Meanwhile, countries that are more responsive to citizen and consumer demands for safety and privacy protection, including European states, may lag behind in AI innovation.

A REALISTIC APPROACH

The AI revolution is a global phenomenon wherein opportunities and challenges are inextricably intertwined, and where technological advances are enmeshed with global trade, investment, and supply chains. In such a context, devising a sustainable and legitimate process for designing, agreeing upon, deploying, and regularly updating global governance is key to managing the rise of AI to benefit society. Global governance is needed to shape the competitive landscape to avoid a race to the bottom that would threaten ethical, safety, and human values while raising standards for beneficial AI innovation. If it is to remain robust and relevant over time, an effective system of AI governance must be anchored in current global governance realities, including the central role of the nation-state and legal norms. It must nevertheless be able to adapt to changing power dynamics, especially the growing influence of transnational actors such as digital multinationals and the rising influence of soft law (codes of conduct and practices, technical standards, and so on).

Likewise, it should also be shaped around an understanding that the rise of AI policy is the result of the complex and dynamic sociotechnical system in which science, technology, and societies influence and even coproduce one other globally and locally (Jasanoff, 2004). New technologies and innovations deployed into societies continuously impact and redefine values and norms, which in turn influence policies and laws both hard and soft. Likewise, the evolution of values, norms, and governance also continuously shape technoscientific developments.

SECTORS AT STAKE

Another key is the ability to understand where and how specific principles and norms—both hard and soft law—should be inserted into existing governance regimes (such as those regulating trade and investment, arms control, human rights, climate change, internet governance, and data regulation) or if, in contrast, new processes, regimes, or institutions need to be designed and deployed.

The stakes at play are high, and power and identities are shifting from nation-states to a more complex transnational, global construct that is increasingly controlled by large technology companies. Given this, it is crucial for diverse stakeholders to participate in AI governance, including government, industry, academia, nonprofits, NGOs, and civil society. Capacity and buy-in would be increased by an inclusive, multistakeholder approach to discovery, debate, and redefinitions around values, ethical principles, the design of international agreements, and their implementation and monitoring. This type of approach would also raise legitimacy and credibility among the public, a key aspect given the current epistemic crisis that many societies are going through, whereby citizens have become skeptical of expert knowledge, as well as of industry or governments' incentives and constraints. The process should be deeply interdisciplinary, reaching beyond science, engineering, and business to actively involve philosophers, artists, sociologists, political scientists, writers, and movie producers. In a world that is increasingly defined by digital mass media and entertainment, these actors play a key, and often underestimated, role in shaping governance through their ability to forge the "collective imaginaries" or narratives which shape public perception and technological evolution pathways.

AN IPCC FOR AI

At this juncture, it would seem to be a priority for there to exist a legitimate process for establishing consensus among stakeholders regarding the nature, dynamics, impacts, and related challenges in the rise of AI. The global coordination challenges around assessing climate changes and its causes are a relevant example. As is the case with the galaxy of factors and drivers contributing to climate change, AI is a complex, pervasive phenomenon that is distributed across society and the economy. It cannot be attributed to a finite set of producers but is instead intertwined with strategic political, trade, and investment interests across states. Accordingly, a potentially relevant model is the Intergovernmental Panel on Climate Change (IPCC).

Under the auspices of the United Nations, the IPCC set a widely acknowledged example of an inclusive, multistakeholder platform for international consensus-building around a matter-of-fact approach to tackling climate change. The IPCC has served as the foundation for designing, implementing, and enforcing global governance and policies, which culminated in the Paris Agreement.

While launching France's national AI strategy on March 29, 2018, President Emmanuel Macron called for more international coordination and proposed the creation of "an IPCC for AI." A large, deeply interdisciplinary group of scientists and experts on AI which performs regular assessments nested in a solid scientific process could play a key role in forging a global consensus on the main challenges to be addressed in AI development.

Likewise, the launch of a global round table on AI governance represents a first step in the process of establishing international, multistakeholder dialogue on this issue. Hosted by the government of the United Arab Emirates and organized by The Future Society at the 2018 World Government Summit in Dubai, the round table kick-started strategic conversations among experts, practitioners, policymakers, and scientists from around the world on the challenges, policy options, and pathways for governing the rise of AI.

COMBINING HARD AND SOFT LAW

Building on this platform, a realistic approach to governance could then draw upon a combination of hard and soft law to raise the bar in safe, ethical AI development. A combination of these two approaches to governance would combine flexible, adaptable governance with legally binding enforcement mechanisms. Applying lessons from successful models of technology governance is a more pragmatic approach than designing new, untested governance regimes. As a starting point, the IEEE technical standards and codes of conduct, and the EU General Data Protection Regulation (GDPR) provide relevant examples of soft and hard governance, respectively. Among other things, these could inform the development of a global governance for AI frameworks.

Soft governance, including industry standards and codes of ethics, offers a promising approach for managing AI innovation. Governance should be flexible and adaptable as technologies and their impact on society evolve at increasing speeds. Regulations that are set in stone often lack the agility to react and adapt effectively to fast-changing technology. Moreover, citizens' values and preferences regarding technology also evolve. The IEEE is a highly relevant example of industry standards aiming to govern safe and ethical autonomous and intelligent systems. The IEEE uses codes of conduct and technical standards to shape technology development. Technical standards can shape AI development from within the industry rather than by relying on external regulation.

Meanwhile, hard governance, including binding legislation, plays a crucial role by creating a level playing field and anchoring technological change in a given value system. A relevant example is the GDPR, which creates a rigorous legal regime applicable to all organizations that collect, store, process, and circulate personal data in Europe. Because of the critical mass of the European digital market with its five hundred million consumers, GDPR may become a global gold standard. Having had to adapt to its demanding rules and regulations to do business in Europe, digital multinationals may seek to standardize operating procedures globally based on it.

WIDE-REACHING SOCIAL BENEFITS

AI will rapidly and significantly transform societies. The rise of AI involves an unprecedented combination of high-stakes, destructive dynamics, including major societal risks and strategic competition among states and firms as part of a global race. In this context, global, multistakeholder governance to coordinate the rise of AI is needed to raise the bar and shape its development to ensure it brings broad societal benefits.

An inclusive international platform modeled on the IPCC could help build consensus and agreement on the key challenges and concerns in AI. Next, a combination of soft and hard governance and policy modeled on successful examples such as the IEEE and Europe's GDPR would represent an incremental and pragmatic approach to governance.

FIGURE 3
FEEDBACK IN THE NEW GLOBAL ORDER

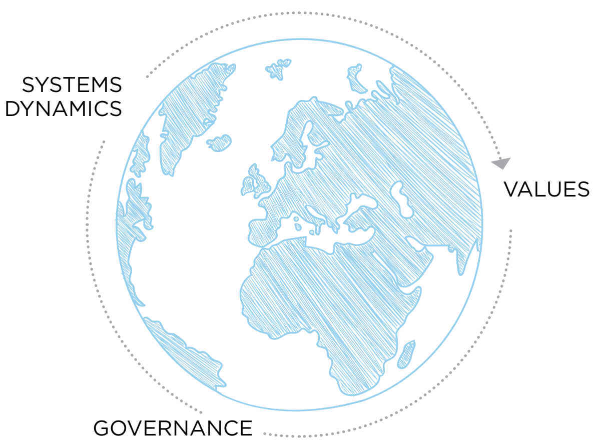

Note: Sociotechnical system dynamics, values, and governance affect and inform each other in a continuous global feedback loop constitutive of an emerging—yet patchy—new global order.
Source: Nicolas Miailhe, The Future Society

NOTES
[1] Interview with Dr. Roman V. Yampolskiy. Simulation, April 3, 2018.

REFERENCES
Armstrong, S., Bostrom, N., and Shulman, C. 2013. "Racing to the Precipice: A Model of Artificial Intelligence Development." FHI Technical Report No. 2013-1. Future of Humanity Institute, Oxford University.
Bostrom, N. 2006. "What is a Singleton?" *Linguistic and Philosophical Investigations* 5 (2), 48–54.
Cerulus, L. 2018. "Macron: France to Invest Nearly Ć.5B for AI until 2022." Politico. March 29.
Chui, M., Manyika, J., Miremadi, M. et al. 2018. "Notes from the AI Frontier: Insights from Hundreds of Use Cases." Discussion Paper, McKinsey Global Institute. April.
Ding, J. 2018. "Deciphering China's AI Dream: The Context, Components, Capabilities, and Consequences of China's Strategy to Lead the World in AI." Report for the Governance of AI Program. Future of Humanity Institute and Oxford University.
European Commission. 2018. "Artificial Intelligence: Commission Outlines a European Approach to Boost Investment and Set Ethical Guidelines." European Commission Press Release. Brussels, Belgium. April 25.
Goldfarb, A., and Trefler, D. 2018. "AI and International Trade." NBER Working Paper No. 24254. Cambridge: NBER.
Jasanoff, S. 2004. "Ordering Knowledge, Ordering Society." In: States of Knowledge: The Co-Production of Science and Social Order." London: Routledge.
McKinsey & Company. 2016. "Car Data: Paving the Way to Value-Creating Mobility: Perspectives on a New Automotive Business Model." McKinsey & Company. March.
Miailhe, N. 2018. "Competing in the Age of Artificial Intelligence: The State-of-the-Art of AI and Interpretation of Complex Data." *Focus (SCOR Global P&C)*. March. 11-14.
Russell, S. 2016. "Q & A: The Future of Artificial Intelligence." Available at: http://people.eecs.berkeley.edu/~russell/temp/q-and-a.html.
Samuel, A.L. 1959. "Some Studies in Machine Learning Using the Game of Checkers." *IBM Journal* 3 (3): 210–229.
The Future Society. 2018. "Governing the Rise of Artificial Intelligence." Available at: http://ai-initiative.org/ai-consultation/.

MAP OF NATIONAL AI STRATEGIES
Countries with local AI plans

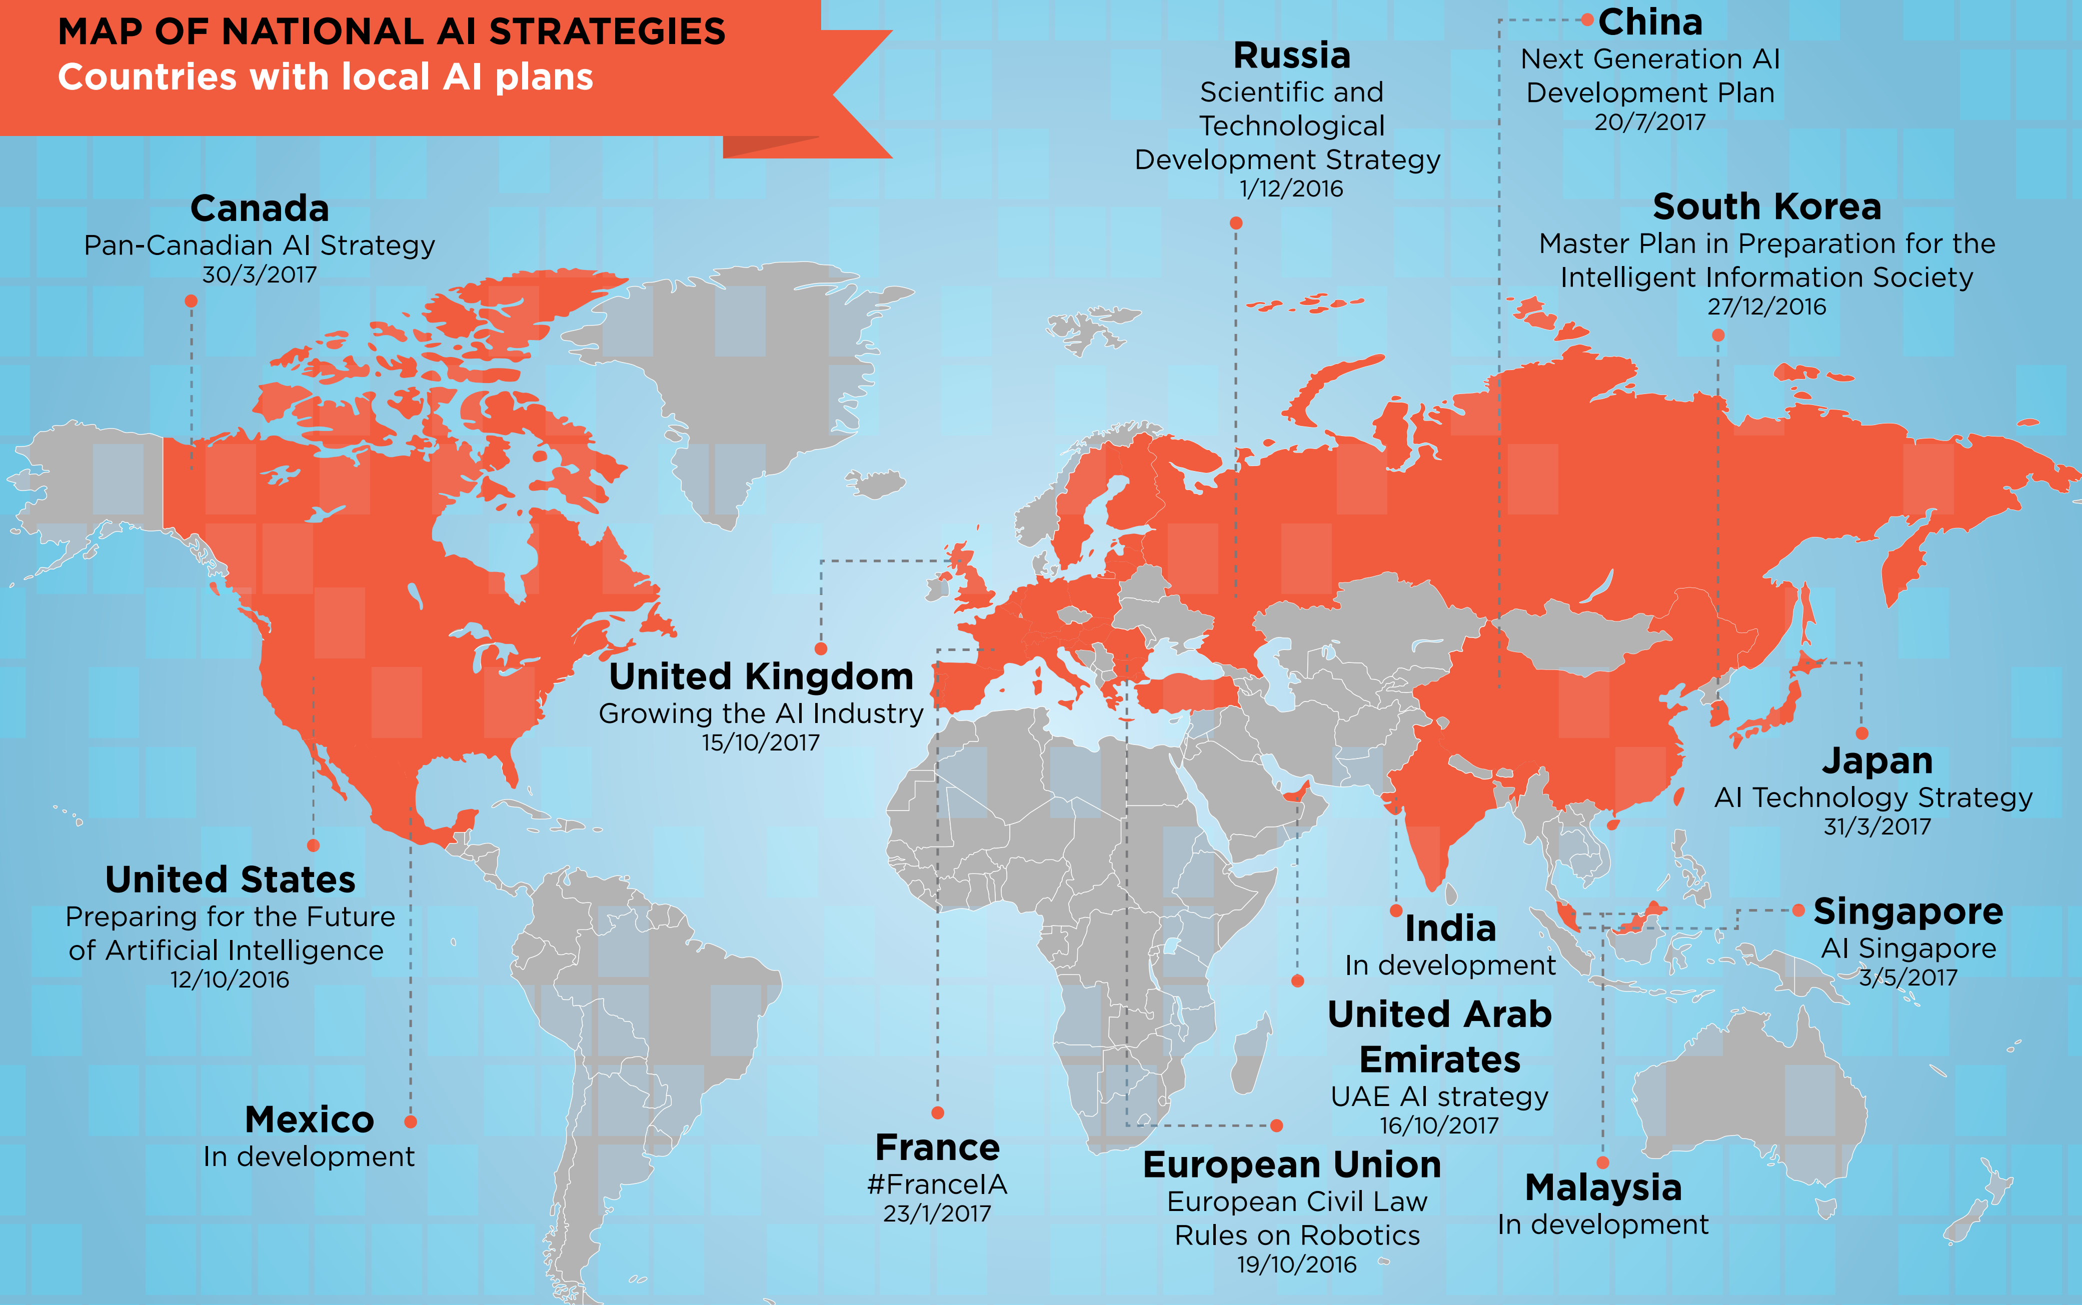

Some countries are focusing their resources on developing AI and AI applications in areas that they consider vital to local development (ICTs in India, the care economy in Japan, energy in China, etc.)

Other countries have published reports on the risks and opportunities that AI may bring without establishing concrete development targets (such as the United States and the European Union).

Geoff Mulgan
University College London
and CEO of Nesta

"Latin America needs to stop giving away its data"

AS CEO OF THE UNITED KINGDOM'S NATIONAL ENDOWMENT FOR SCIENCE, TECHNOLOGY AND THE ARTS (NESTA), GEOFF MULGAN IS A LEADING GLOBAL FIGURE IN THE WORLD OF NEW TECHNOLOGIES AND INNOVATION. IN THIS REPORT, HE SUGGESTS THAT THE COUNTRIES OF LATIN AMERICA NEED TO WORK TOGETHER ON PROJECTS FOR ARTIFICIAL INTELLIGENCE (AI) APPLICATIONS IN SECURITY, HEALTH, AND EDUCATION AND TO SEEK EFFICIENCY THROUGH THE BENEFITS THAT COME FROM OPERATING ON A LARGER SCALE. HE ALSO RECOMMENDS THAT THE REGION TAKE ADVANTAGE OF EXISTING DEVELOPMENTS AND THE INVESTMENTS IN R&D THAT HAVE BEEN MADE IN OTHER COUNTRIES TO SOLVE LOCAL PROBLEMS.

Are people's expectations around AI well-founded?

There is a lot of excitement around AI, a great deal of hype, and some misunderstandings around what AI can actually do. Some ideas come more from Hollywood movies rather than from reality. We need a healthy dose of skepticism before starting a conversation about AI. Despite that, I think that over the next 20 to 30 years, various kinds of AI will change pretty much everything governments do. AI is already transforming everyday repetitive processes such as tax collection or handling files or applications. Predictive algorithms have been already used for some time now to predict, for example, who is more likely to go to a hospital, or which prisoners will reoffend, or which students will perform best at school. This kind of AI will be used more and more, and it will become increasingly accurate at predicting behavior.

How should public policies adapt to these changes?

We have seen rapid transformations in the interfaces between governments and citizens using chatbots of all kinds to handle conversations and answer questions. Indeed, I predict that in a few years, chatbots will be our main form of interaction with the health system, rather the front office of a hospital or a doctor's surgery. There are also a lot of uses for these technologies in education, such as the personalization of maths teaching. There are even some democracy-related uses, to enhance public consultations and make it easier to understands different citizens' opinions.

How can governments take maximum advantage of the potential of AI?

In most governments, there certainly is a lack of basic talent and skills on how to use AI effectively. The other missing part everywhere is good test beds. We don't need to reinvent the wheel all the time: the task of test beds is to try out these tools and adapt them to make them more useful. The lack of test beds is true the world over. They're not very difficult or expensive to organize, whether for transportation planning, or the education system, or social security. You don't need basic research, because companies like Amazon or Google and countries like China are already investing massively in that. But we can do simpler things, including adapting what is already there in the cloud to solve everyday problems. I don't know if there are Latin American institutions that are consciously dealing with that.

GEOFF MULGAN RECEIVED HIS PH.D. IN TELECOMMUNICATIONS FROM THE UNIVERSITY OF WESTMINSTER. SINCE THEN HE HAS BECOME A KEY FIGURE IN INNOVATION IN THE UNITED KINGDOM AND HAS PUBLISHED A SERIES OF BOOKS THAT ARE ESSENTIAL READING ON THIS ISSUE. THEY INCLUDE:

- *COMMUNICATION AND CONTROL: NETWORKS AND THE NEW ECONOMIES OF COMMUNICATION*. NEW YORK: GUILFORD PRESS. 1991.
- *POLITICS IN AN ANTIPOLITICAL AGE*. CAMBRIDGE: POLITY PRESS. 1994.
- *CONNEXITY*: HOW TO LIVE IN A CONNECTED WORLD. BOSTON: HARVARD BUSINESS SCHOOL PRESS. 1997.
- *GOOD AND BAD POWER: THE IDEALS AND BETRAYALS OF GOVERNMENT*. LONDON: PENGUIN. 2006.
- *THE ART OF PUBLIC STRATEGY: MOBILIZING POWER AND KNOWLEDGE FOR THE COMMON GOOD*. OXFORD: OXFORD UNIVERSITY PRESS. 2009.
- *THE LOCUST AND THE BEE: PREDATORS AND CREATORS IN CAPITALISM'S FUTURE*. PRINCETON: PRINCETON UNIVERSITY PRESS. 2013.

What risks does the widespread use of this technology bring?

We already can see a huge number of uses associated with risks, which come from the data underlying AI, which can be very biased or distorted in different ways. There are a lot of big questions about who should own the data, who should own the algorithms, and who should be accountable when things go wrong. Probably the biggest challenge in this matter around government is skilling yourself up, becoming a smart discerning customer who isn't vulnerable to overselling or overhype, but without going to the opposite extreme and seeing everything through a negative lens. We need to address the paradox of how AI—which has emerged from public investment in the military, intelligence services, or universities—is still so behind in meeting public needs like improving the social security system, transportation, or schooling. This is an opportunity for organizations like the IDB to play a strategic role by identifying major needs and pulling together government investments in experiments to test AI practically.

Do you think that automation will lead to mass unemployment?

We have carried out a major study about the future of jobs in the UK and the US. The conclusions for those countries focus on the skills that will be needed in jobs and they suggest that the impact of automatization will be much lower than what the media are currently reporting. This is partly because so many jobs combine elements that can't be automated. Indeed, we have concluded that some public-sector jobs like teachers and doctors will probably grow in number as automatization will produce more productivity in other parts of the economy.

What will the dynamics be like in developing countries?

The situation isn't as clear. Probably every government needs to ensure that, at the very least, they are accelerating the reskilling process so that people can adapt to a very unpredictable labor market environment. In some countries, we have been advocating adult learning, for example, on the assumption that whatever happens, people will need to adjust their career or their jobs more often throughout their life. They will need help to navigate through that, assess their own skills, and understand what extra skills will make them more resilient against future effects of automation in the labor market. And paradoxically this is a field where using AI and data tools can empower people to better thrive through the turbulence that AI is causing. There is a relatively easy action path for governments and organizations to make the labor market smarter and not just see these tools through the lens of fear.

LATIN AMERICA HAS ENORMOUS INTELLECTUAL CAPITAL THAT IT COULD USE TO SOLVE LOCAL PROBLEMS

How can Latin America make best use of this disruptive technology?

I'm on the steering committee of an interesting group about education in Latin America called Suma, which the IDB is also involved in. I see education as a good field for leaping up in terms of collective intelligence at the continent level. Where the whole project could go would be much more about drawing on data around what is working and not working in schools, mobilizing teachers to run low-level experiments in different methods of teaching, or different ways of using technology. Another idea is creating content while test beds try out new solutions like using AI for teaching maths. All of that could be organized much more effectively at a regional scale, involving the whole of Latin America. We need to think, learn, experiment together, rather than saying that those are things that only happen in universities, or labs, or in the offices of isolated institutions. Integration and collaboration are key.

How can AI help solve the region's main problems?

Take for instance gun violence: we already have tools that can predict where gun violence might happen. We can also use it to pool collective intelligence from police officers and citizens to take action, such as by providing courses for retraining people who have left gangs. All of that is possible, but of course these kinds of applications remind us that when we use collective intelligence there will also be some people that will try to undermine it, distort the information, and skew behavior. It is fantastic to use these tools to pool information and data but you have to be constantly building up the immune system to fight the enemies of collective intelligence who will try to disrupt the good work of a society. For example, when we introduced open data to policing here in the UK, it was obvious that this would directly benefit criminals because they could see which places had the worst policing.

How would you suggest we tackle the issue of data privacy?

We are entering a phase of geopolitical competition around AI. There is huge investment in China, the US, and, to a lesser extent, in Europe to be at the forefront of this technological revolution. Regarding Latin America, I would say that the key is ensuring you don't give your data away to fuel technology that is controlled by others. The region needs to be smart about this environment. Data is not the new oil, that's a misleading metaphor, but it definitely is a huge resource of value and most AIs can't run without data. Up to now, Latin America has certainly been giving its data away, mainly to foreign companies. You need to work together as a region. If you can't collaborate on that, then there is no chance of putting your best foot forward in this radically different era in human economic history. Latin America has enormous intellectual capital. The question is how to harness this intellectual capital to solve real problems.

CASE STUDY

India's Aadhaar Experiment

Samir Saran and Madhulika Srikumar
Observer Research Foundation, New Delhi

The history of technology in developing economies is replete with failures—lack of access to infrastructure and a dearth of technical capacity have contributed to social resistance to them over time. Aadhaar, India's bold experiment to return agency to its people, however, is poised to become an exception. This national biometric ID project has enabled the digital onboarding of over a billion people by providing citizens with a unique identification number, many for the first time.

For decades, people in India have struggled to officially prove their identity. Until last year, only 5.5% of the population had passports, and while other forms of identity such as the PAN card or voter ID were widely in use, these were susceptible to easy duplication. In India, citizenship is often considered an assumption rather than as an established fact. This problem is not exclusive to India. An estimated 1.1 billion people worldwide do not have identification and a staggering 2 billion do not have a bank account or access to a financial institution via a mobile phone or any other device.

Proving identity at scale is crucial. At the heart of access lies the imperative to provide an identification system that is unique, authentic, reliable, and digital. More importantly, proof of identity is necessary to exercise a wide range of rights. Aadhaar—which literally translates as "foundation"—set out to achieve exactly this, to not just verify citizens' identities but to also provide a foundation for people to interact with the state.

Through biometrics (fingerprint and iris recognition), the Aadhaar card can verify and authenticate a person and enable the delivery of welfare services in an efficient and transparent manner. Aadhaar is the world's largest biometric ID project now that reportedly 99% of Indians aged 18 and above have an Aadhaar number. The unique ID project has created a one-of-a-kind public data infrastructure that rivals the biggest service providers in the world: the US's FANG—Facebook, Amazon, Apple, Netflix—and China's BAT—Baidu, Alibaba, and Tencent. It is the only digital platform in the world that has over 1 billion people on it but is not privately owned. Over 20 countries have expressed an interest in implementing a similar digital identity system. Digital proof of identity is increasingly being seen as a tool for poverty alleviation and the delivery of essential services, and it could become India's soft power export on the global stage.

The Aadhaar identity system uses a twelve-digit number that can be obtained by any resident of India on submitting demographic and biometric data. The program has been operational since 2009. The Unique Identification Authority of India (UIDAI), a statutory body established under the Ministry of Electronics and Information Technology (MeitY), is responsible for issuing identity numbers and collecting data. The national unique identity program, however, only received legislative support in 2016 with the passing of the Aadhaar (Targeted Delivery of Financial and Other Subsidies, Benefits and Services) Act in the Lok Sabha or Lower House.

While the unique identity system is a powerful enabler, bringing citizens into the fold of the formal economy, Aadhaar can also lock people out of the system—elderly and disabled citizens, for instance, struggle to enroll for an Aadhaar ID due to their fading fingerprints.

As with the introduction of any new technology, people and processes form an important part of the value chain to protect users and data. Aadhaar going through a teething period, one that includes instances of unauthorized access leading to illegal storage and use of citizen data. The Indian Supreme Court is currently determining the constitutionality of the program. India's first data protection law is also in the works after a series of public consultation. Some of the country's sharpest minds in industry and government will explore how anonymous public data can be used for public good through machine learning applications and will examine the role of AI in securing the Aadhaar platform.

The Aadhaar digital identity system serves as a unique "financial address," enabling the transfer of benefits directly to beneficiaries' accounts through real-time authentication. Aadhaar-enabled payments (AEPS) is unique to India. This service enables people to make transactions seamlessly merely through their Aadhaar numbers.

The Indian government is promoting financial inclusion by connecting Jan Dhan or newly opened bank accounts, the Aadhaar platform, and mobile numbers (a system known as the "JAM trinity") to create a common financial, economic, and digital space. The Aadhaar digital payments infrastructure has been pivotal in enabling the state to transfer benefits to people—it is reported that 900 million are currently receiving funds through these platforms. In 2017, two welfare schemes made increasing numbers of transfers through the Aadhaar Payment Bridge System (APBS). In May 2017, over 95% of direct cash transfers for grain subsidies and 82% for LPG were carried out through the APBS.

Of the 600 million people in India who own phones, 300 million use smartphones. However, due to lack of last mile connectivity, a significant section of the Indian population has remained outside the financial system. Banking the unbanked is a huge logistical undertaking for the government, necessitating the opening of countless physical bank branches. Indian policymakers have astutely circumvented this by adopting contextual technological solutions instead. The government has leveraged the growing penetration of mobile phones in rural India to deliver banking and other services directly to a demographic that has largely been inaccessible for the last several decades. Aadhaar, by providing digital solutions, has overcome the challenges of distance, cost, and literacy.

The direct benefits transfer system has removed the middlemen at the state, municipal, and neighborhood level, reportedly plugging leakages as high as USD500 million. This has brought down the overall cost of transactions. The digital identity scheme has been critical in eliminating duplicate and fake identities. Jan Dhan has enabled the opening of 300 million bank accounts since its launch. It is reported that the number of zero-balance accounts dropped from 60% in 2015 to 23% in 2017.

To promote digital payments, the government is changing the way it operates by seeking to make every payment over INR5000 only using electronic means. The data shows that the government is receiving over 95% of tax receipts electronically, including income, excise, and custom taxes. The state is now taking strides to receive nontariff revenue through digital payments.

It has also implemented various policy initiatives to incentivize digital payments. The transaction costs of intermediaries have been absorbed by the government through their own technologies to incentivize receivers, banks, and users through reimbursement. The government has created a differential income tax rate for small businesses which are transitioning from cash to digital payments, reducing the rate from 8% to 6%.

Digital wallets constitute the biggest mode of payment in India outside of plastic cards and have shown 300% growth since the demonetization initiative of 2016. Payments via UPI, USSD, and AEPS have grown by over 1,000%.

By providing authentication services, the Aadhaar database has also spurred the creation of an ecosystem of applications that cater to peoples' everyday needs, which is called the IndiaStack. The IndiaStack, a suite of APIs, provides organizations with a toolkit for delivering services in a presenceless, paperless, and cashless manner. A person can now sign up for a service using just her mobile phone with Aadhaar eKYC, store and verify documents online through DigiLocker, and pay for services through the Unified Payments Interface (UPI), India's interoperable public digital payments infrastructure. In addition to building layers to deliver these services transparently, IndiaStack has also incorporated individual rights by design through a "consent" layer, which will enable users to control their information and how it is shared.

The Aadhaar ecosystem has enabled private corporations to build on this layer. Private banks in India have adopted the Aadhaar framework to verify customer identities when they open new bank accounts. The integration of WhatsApp with the government's UPI is also symbolic, representing a convergence between one of the jewels in Silicon Valley's crown and a locally created architecture.

One of the most important benefits of Aadhaar has been the low cost of enrollment, which is said to be approximately USD1.2 dollars for providing each individual with a unique identity and bringing them into the fold of the formal economy. The true cost, however, lies in the long-term investments that have been made in state institutions to safeguard individual rights.

Aadhaar, besides powering India's transition to a digital economy, has also created pathways for participatory governance. The government e-marketplace (GEM) initiative facilitates online procurement of essential goods and services for government departments, promoting entrepreneurship. Anyone with a phone can gain direct access to government departments and bid for projects in a transparent manner, authenticated and verified by the Aadhaar database. This technology-driven governance has also seeped into other government services. Through the mygov.in portal, people can give their opinions on government programs, vote for alternatives, and provide inputs on policy design. Ministries now frequently provide draft policies for consultation on this platform. This symbolizes a move towards participatory governance in India.

The digital revolution presents an opportunity not just to transfer technologies but also to transfer the normative institutions that safeguard rights and open spaces. India and emerging economies in Latin America are in a position to create innovations that will improve the lives of the next six billion people.

The **Moral Risk** in *Machine Learning*

Juan Miguel Lavista
Microsoft

DISCRIMINATION, FALSE PREDICTIONS, SYSTEMS THAT ARE EASY TO CHEAT. THE AUTOMATION OF PRODUCTION POSES NEW ETHICAL DILEMMAS. HOW CAN WE GUARANTEE THAT THE ONLY PARTS OF HUMAN BEHAVIOR THAT ALGORITHMS LEARN ARE THOSE WORTH IMITATING?

It was the summer of 2016, and as most farmers in India began to sow their groundnut crops, a small group of 174 farmers held out, waiting on a text that would revolutionize the farming industry.[1]

For centuries, farmers in Andhra Pradesh and Karnataka region in India had been using an ancient method to forecast the optimal sowing date for their groundnut crop. This technique suggested sowing in early June to take advantage of the monsoon season. But recent changes to the weather patterns resulted in unpredictable monsoons, causing poor crop yields.

Chinnavenkateswarlu, a farmer from Bairavanikunta village, was among the 174 farmers that were waiting for the text message that would alert them of the best sowing date.

Finally, on October 28, Chinnavenkateswarlu received a text message on his phone instructing him to begin the sowing process. The text message was the output of the artificial intelligence (AI) Sowing App that was developed by Microsoft in collaboration with the International Crop Research Institute for the Semi-Arid Tropics (ICRISAT). ICRISAT is a nonprofit, nonpolitical organization that conducts agricultural research for development in Asia and sub-Saharan Africa.

To determine the optimal sowing period the system uses AI and machine-learning models that are built using historical climate data, including the season's recorded daily rainfall and the weather forecast for the region.

For Chinnavenkateswarlu, using the system paid off: along with the other 174 farmers, he achieved an average 30% higher yield per hectare.

We can no longer consider AI as a niche, just another part of the distant future. This example is just one of the many cases showing the impact that AI is currently having around the world. Every month two billion users around the globe use search engines like Bing, Google, or Baidu to make more than 180 billion searches. The search results are provided by AI algorithms. The posts and news stories offered to more than two billion Facebook users worldwide are chosen by AI algorithms that use individuals' actions to learn about their interests.

From logistics to medicine, from individualized interest rates when applying for a credit card to Skype's automatic translation, which allows users around the world to communicate, our lives and decisions are shaped by AI algorithms.

HOW DID WE END UP HERE?

The majority of what are considered AI systems today are based on machine-learning algorithms. In conventional programming, humans translate their knowledge into a computer language code. In machine learning, on the other hand, there are algorithms that will use data to learn from it and make predictions.

The power of machine learning is that if we have the right data and algorithms, not only can we solve problems that would previously have been too complex for humans to solve using conventional programming, we can also be more efficient.

In the 1990s there were multiple software solutions for handwriting recognition. These products involved huge teams and months of software development.[2] Today, using the MNIST database[3] (a large database of handwritten digital images), a data scientist can train a convolutional neural network that can achieve more accurate results in less than 100 lines of code.

Machine learning is not a new field, in fact, a lot of the algorithms used today were created 20 to 40 years ago.[4] So, we need to ask ourselves, why now? As I mentioned before, even though algorithms are needed for machine learning, another key component is data and the ability to process it.

The remarkable growth of the internet and the online world, combined with the mass adoption of smartphones, has created huge amounts of data, a key requirement for algorithms to be able to "learn." Likewise, the cost of storing and processing this data has dramatically decreased (see figure 2).

Computer processing power and storage capacity per unit cost have been following Moore's law. Today, a US$3,000 GPU computer has more processing power than the US$500 million supercomputer that NEC built in 2001 and called the "earth simulator," which at that point was considered the most powerful computer in the world. Ten terabytes of data, equivalent to the size of the US Library of Congress's

FIGURE 1
SAMPLE IMAGES FROM MNIST TEST DATASET

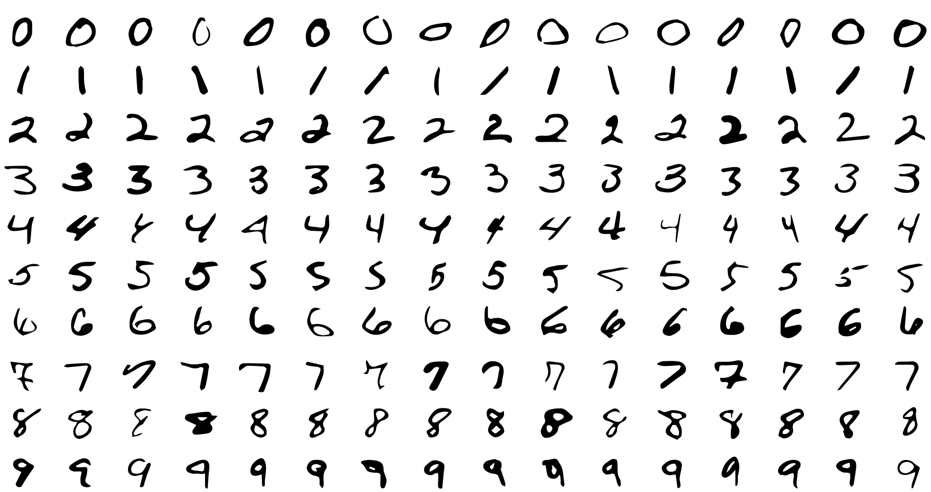

Source: Josef Steppan.

print collection (26 million books) costs US$180 today but would have cost over US$200,000 in the year 2000 (table 2).

This phenomenon has resulted in a vast number of opportunities for using machine learning to solve problems. Today, the main competitive advantage of companies working in this space is no longer software or algorithms. In fact, most companies rely on the same open source solutions and algorithms to do AI. Instead, the real value and differentiation factor is data.

Being able to learn and generalize from data provides amazing power. Every day, scientists around the world are finding new ways to solve problems using data, AI, and machine learning. Nevertheless, it's important to consider that our solutions will only be as good as the data. If there are issues or biases in the data, the output of the algorithms will be affected. The following section contains a few lessons that need to be considered when solving problems using AI or machine learning.

NUMBERS CAN BE DECEIVING

In 1991, Dr. Diane F. Halpern of California State University at San Bernardino and Dr. Stanley Coren of the University of British Columbia published a paper with an alarming conclusion. The researchers took a random sample of individuals that died and asked their family members if they were left-handed or right-handed. The finding was disturbing: left-handed people were dying nine years earlier than right-handed ones (Coren and Halpern, 1991).

This paper was published in the *New England Journal of Medicine*, which is one of the most prestigious medical journals in the world. If the findings were correct, this would mean that being left-handed was as bad as smoking 120 cigarettes per day (Barnes, 2013).

The problem with this study was that the researchers did not take into account the fact that, for a long time in history, being left-handed was perceived as something bad so parents forced their children to be right-handed. Eventually, parents stopped doing this, and it generated an artificial increase in the left-handed population. This artificial increase is responsible for giving the illusion that left-handed people die younger.

The issue is that if a life insurance company uses machine learning and left-handedness is one of the attributes it bases its calculations on, then an AI/machine-learning model will use this information and will wrongly predict that left-handed people die younger, potentially resulting in higher premiums for left-handed people.

What can we learn from this? There is a bias to most of the data we collect. It is fundamental to understand this bias and its potential effects on the machine-learning models.

ALGORITHMS AND DISCRIMINATION

In the early 1980s, employees from St George's Hospital Medical School in London, England, decided to use an algorithm to automate the first round of its admissions process. The algorithm was built using historical data from previous applications and produced results that were 90%–95% similar to what a human panel would have decided (Collier and Burke, 1986).

In 1986, it was proven that the algorithms provided much lower scores to women and those from racial minorities, thereby reducing their chances of being interviewed. The problem was not that the algorithm was introducing a new gender or racial bias but rather that the algorithm was learning this bias from historical data and perpetuating it (Lowry and Macpherson, 1988).

As Dr. Cathy O'Neil (2017) clearly states, the benefit of algorithms is that we can easily interrogate them and check for these biases. We can test algorithms to understand what effect gender or race has on the output while controlling for all other factors, and we can use this information to remove the gender/race bias.

What can we learn from this? AI/machine-learning models learn from data. This means that if we train the algorithms using data that is in some way discriminatory, the algorithms will also learn to discriminate. It is easy to test algorithms and understand whether or not they are acting in a discriminatory fashion.

CORRELATION DOES NOT IMPLY CAUSALITY

In the US, people that drive a Mercedes live longer than the general population. This correlation is not a spurious correlation: it can be explained by the fact that Mercedes owners will on average have significantly higher incomes than the general public, and a set of confounding variables related to higher income can explain why they live longer, on average (Dickman, Himmelstein, and Woolhandler, 2017).

Of course, if people that cannot afford a Mercedes try buying one to increase their longevity, it will perversely result in the opposite outcome by squeezing their limited budget. This example seems obvious, but unfortunately most people do not understand it. Gallup did a survey a few years ago asking a straightforward question: "Do you believe correlation implies causation?" Surprisingly, 64% of Americans said that they did (Sobel and Shiraev, 2016).

The problem of confusing correlation with causation is so prevalent that almost every day, articles are published in reputable media that clearly confuse the two. A few months ago, a paper was published in the *Journal of the American Medical Association* titled "Inequalities in Life Expectancy Among US Counties, 1980 to 2014" (Dwyer-Lindgren et al., 2017). The paper was all over the news, and when the *Miami Herald* reported, "Want to live longer than the 80.9 years of your life expectancy in Miami? Move to Colorado" (Robertson, 2017), it was clearly confusing correlation with causation.

What can we learn from this? We need to remember that correlation does not imply causation. Predictive models don't require the data they use to be based on causal relationships. When presenting the results from models, it is critical to make this point clear as a significant portion of the population might treat connections as being causal.

FIGURE 2
COST OF HARD DISK PER GIGABYTE (US$)

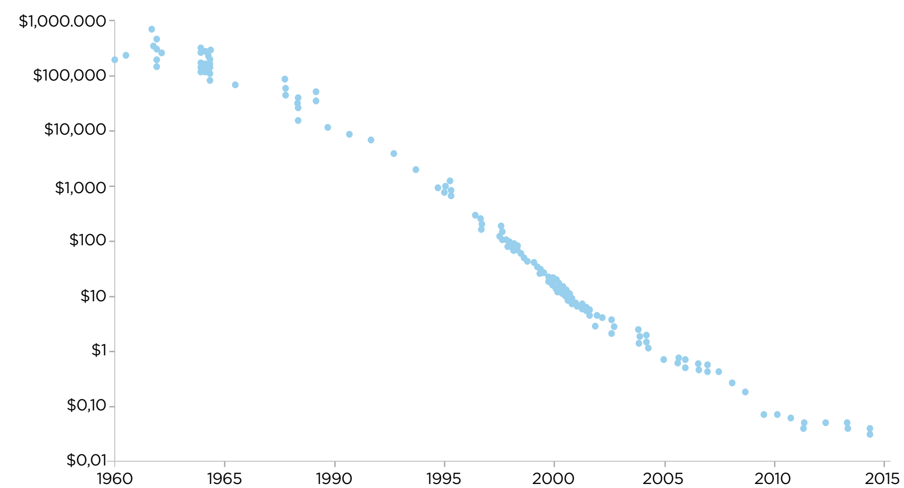

Source: mkomo.com.

OBJECTIVES AND MEASUREMENTS

During the 1800s, the British colonial government in India was concerned about the high numbers of people bitten by cobras in Delhi (Dubner, 2012). To deal with this problem, the government made a policy decision to pay a bounty for every dead cobra brought forward. The policy worked very well for the first few months: cobra killing increased, and, as a result, the number of cobras in the streets declined.

However, a few months after the policy was introduced, something strange started to happen. The number of dead cobras continued to increase, but for the first time, the number of people bitten by cobras also increased. Killing cobras became a business, so people started to breed cobras for the income. Not only was the public policy not working, it was actually exacerbating the problem it was intended to solve.

In the 1970s, Charles Goodhart, a former advisor to the Bank of England and Professor Emeritus at the London School of Economics, described this exact problem: "When a measure becomes a target, it ceases to be a good measure" (Strathern, 1997).

Machine-learning models require that the set of signals/features used in them have information with predictive power. However, the relationship between these signals/features and the outcome do not necessarily need to be causal. This means that the feature might be indicative/correlated but not necessarily the cause of what we are trying to predict.

For example, let's say we need to predict [C], but we can only measure [B] as a feature, and [B] does not affect [C]. The real cause that affects [C] is [A], but we cannot measure [A]. [A], on the other hand, also affects [B], so we can use [B] as a way to predict [C].

Let's say we estimate this model and

TABLE 1
COMPARISON OF GPU AND SUPERCOMPUTERS BY COMPUTING POWER AND Cost

	NVIDIA TITAN V	NEC EARTH SIMULATOR
PERFORMANCE	110 TERAFLOPS	41 TERAFLOPS
PRICE	US$ 2,999	US$ 530 MILLION
YEAR OF MANUFACTURE	2017	2001

Source: Compiled by the author.

it gives us very good predictions for [C]. However, we later release the model to the public, who takes advantage of this information to start targeting changes directly in [B]. At this point, the information provided by [B] to the model is no longer associated only with [A], so it loses power for predicting [C].

For example, Pagerank provides a way to rank the importance of a website based on which other websites have links to this particular website (Page et al., 1999). The basis of the model is that if a document is important or relevant, other websites will reference it and include links to it, and when the links occur naturally, the ranking works. However, once this relationship becomes public knowledge, there are clear incentives to play the system (as was the case with the cobras in Delhi). For example, users might pay others to link to their site in order to increase its search ranking. However, if we merely artificially increase the links to a website it will not make the website more relevant.

Credit scoring is another example of this. The design objective of the FICO credit risk score is to predict the likelihood that a consumer will go 90 days past due (or worse) on their payments in the 24 months after the score is calculated. Credit scores are another place where disclosing rules end up damaging the model. For example, myFICO (2018) states: "Research shows that opening several credit accounts in a short period of time represents a greater risk—especially for people who don't have a long credit history." This type of study shows correlation but not causation. These findings can help predict credit risk; however, by disclosing these rules, the credit agency incurs the risk that users will learn how to play the system and therefore damage its credit risk predicting power.

What can we learn from this? If the relationship between the feature and the outcome is not causal, especially if the signal/feature is easy to change—for example, by buying links in the example above—and there are reasons why people have incentives to affect the actual outcome, then there may be a risk of users playing the system. It's important to understand and evaluate risks, as well as to monitor systems periodically.

FIGURE 3
CAUSALITY AND CORRELATION

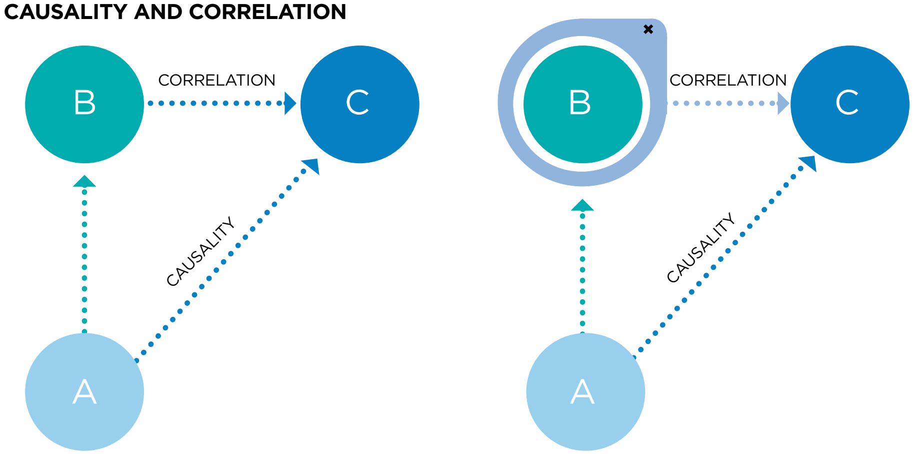

Source: Compiled by the author.

TABLE 2
COST OF STORING LIBRARY OF CONGRESS BOOKS

YEAR	COST
2001	US$ 200,000
2017	US$ 180

Source: Lyman and Varian (2000).

ETHICAL LEARNING

AI and machine learning are already improving our lives today and will change the world tomorrow in ways that are unimaginable to us now. The power of AI relies on data, and we expect that the trend of cheaper storage and processing power will continue, and this will provide even more opportunities for problems to be solved using data.

At the same time, it is critical for us to understand that learning from data has risks that must be considered. If we are learning from human behavior, we need to understand that AI/machine-learning models can learn good and bad things, including the ability to discriminate against people or groups. At the same time, AI/machine-learning models have the power to show that there is discrimination in the first place and can provide a path to fixing it.

NOTES
[1] For more on this topic, see Microsoft (n/d).
[2] Wikipedia, s.v., "Handwriting Recognition." Last modified November 30, 2017, https://en.wikipedia.org/wiki/Handwriting_recognition.
[3] LeCun, Y., Cortes, C., and Burges, C.J.C. The MNIST Database.
[4] Wikipedia, s.v. "Timeline of Machine Learning," last modified February 17, 2018, https://en.wikipedia.org/wiki/Timeline_of_machine_learning.

REFERENCES
Barnes, H. 2013. "Do Left-Handed People Really Die Young?" BBC News. September 7.
Collier, J., and Burke, A. 1986. "Racial and Sexual Discrimination in the Selection of Students for London Medical Schools." *Medical Education* 20 (2): 86–90.
Coren, S., and Halpern, D.F. 1991. "Left-Handedness: A Marker for Decreased Survival Fitness." *Psychological Bulletin* 109 (1): 90–106.
Dickman, S., Himmelstein, D.U., and Woolhandler, S. 2017. "Inequality and the Health-Care System in the USA." *The Lancet* 389 (10077): 1431–1441.
Dubner, S.J. 2012. "The Cobra Effect: A New Freakonomics Radio Podcast." *Freakonomics.* October 11.
Dwyer-Lindgren, L., Bertozzi-Villa, A., Stubbs, R.W. *et al.* 2017. "Inequalities in Life Expectancy Among US Counties, 1980 to 2014. Temporal Trends and Key Drivers." *JAMA Internal Medicine* 177 (7): 1003–1011.
Lowry, S., and Macpherson, G. 1988. "A Blot on the Profession." *British Medical Journal* 296 (6623): 657–658.
Lyman P., and Varian H.R. 2000. *How Much Information?* Available at: http://groups.ischool.berkeley.edu/archive/how-much-info-2003/
Microsoft. s/f. "Digital Agriculture: Farmers in India Are Using AI to Increase Crop Yields." Microsoft News Center India.
MyFICO. 2018. "What's in My FICO Scores." Available at: https://www.myfico.com/credit-education/whats-in-your-credit-score/.
O'Neil, C. 2017. "The Era of Blind Faith in Big Data Must End." TED Talk. April.
Page, L., Brin, S., Motwani, R. et al. 1999. "The PageRank Citation Ranking: Bringing Order to the Web." Technical report. Stanford InfoLab.
Robertson, L. 2017. "Want to Live Longer than the 80.9 Years of Your Life Expectancy in Miami? Move to Colorado." *Miami Herald* May 17.
Sobel, R., and Shiraev, E.B. 2016. *People and Their Opinions: Thinking Critically About Public Opinions.* New York: Routledge.
Strathern, M. 1997. "'Improving Ratings': Audit in the British University system." *European Review*. 5 (3): 305–321.

CASE STUDY

AI Procurement

Criteria for Government Procurement

Doaa Abu Elyounes
Harvard University

There is a growing understanding that governments can use public procurement to boost artificial intelligence (AI). In the United Kingdom, for example, the House of Lords Select Committee on AI published a report in April 2018 that urged the UK government to take action and smartly invest in AI to boost the economy, help solve social problems, and save on the resources allocated annually to goods and services across the government.

Among other benefits of smart government investment in AI, the report estimated that the use of AI virtual agents across the government could save up to GBP4 billion per year. In addition, the use of AI would transform the way government agencies work and enable more informed public policy decisions. However, the report also acknowledges that the fact that big companies are developing and operating with huge amounts of data may entail risks that the government should control and maybe regulate. As a consequence, any purchase of an AI tool by the government needs to support the use of AI for social good and in particular to prohibit the use of opaque technology.

One very interesting initiative that the report highlighted is the establishment of a new public sector unit based within the government that is called The GovTech Catalyst. The unit will provide a direct access point to the government for businesses, and it will support public bodies in procuring innovative products.

Other ways of investing in AI include allocating government funds to support AI research, including social science research that assesses not only the technical innovations themselves but also the impact of the expected product on society. A major layer of complexity that arises in this context is that the private companies that the government contracts usually include nondisclosure agreements which prevent access to proprietary codes. Companies justify this request by claiming that intellectual property laws in general, and trade secrets in particular, are needed to protect the code from competitors and from alleged criminals who might modify the way the code functions and otherwise circumvent the technology. It is thus essential to develop a regulatory framework that makes this kind of information exchange more transparent.

Another interesting example is described in a study conducted by researchers at University College London, who developed a machine learning algorithm that can predict the outcome of cases heard by the European Court of Human Rights (Wakefield, 2016). The researchers analyzed hundreds of cases heard by the court and used language processing and machine learning to identify patterns driving judicial decisions (Aletras et al., 2016). The algorithm scanned the cases and looked for patterns that would help in the final classification of violation or nonviolation of a specific article of the Convention on Human Rights. The model can predict case outcomes with 79% accuracy.

In the United States and in Europe, many police departments in major cities are already using or planning to adopt the use of body cameras in the near future. The contracts that government agencies sign with the companies that develop the tools can ensure that basic values that are important to the government our embodied within the product. These values can include privacy, security, accountability, and transparency.

Policymakers are faced with more and more offers to purchase advanced algorithms. The decision to favor one group of factors over the other is a policy choice and it will affect the results that the algorithm will produce. In machine learning, usually more data equals better prediction. When data is collected for as many risk factors as possible, then statistical analyses can reveal the combination of factors that, when evaluated together, are the most accurate predictors.

Building any AI-based algorithm requires not only the expertise of statisticians but also access to a high-quality database. However, there is no protocol for collecting and maintaining high-quality data and this is a serious hurdle. Since machine learning algorithms are based on training and validation, the quality of the data will certainly be reflected in the results. Another risk is that the training data that we will be using is discriminatory, in which case, the results will also be discriminatory and an algorithm that supposed to be color blind will reinforce biases in the name of science.

Opponents to the use of algorithmic risk assessment tools claim that black box tools produce outcomes that may clash with requirements for transparency. However, algorithms merely assist decision-makers, rather than replacing them completely. In any case, transparency will be an easier target to reach if government organizations insist on exercising their preferences during negotiations with private AI service providers.

REFERENCES

Aletras, N., Tsarapatsanis, D., Preotiuc-Pietro, D., et al. 2016. "Predicting Judicial Decisions of the European Court of Human Rights: A Natural Language Processing Perspective." *Peer J Computer Science* 2:e93.

Wakefield, J. 2016. "AI Predicts Outcome of Human Rights Cases." *BBC News.* October 23.

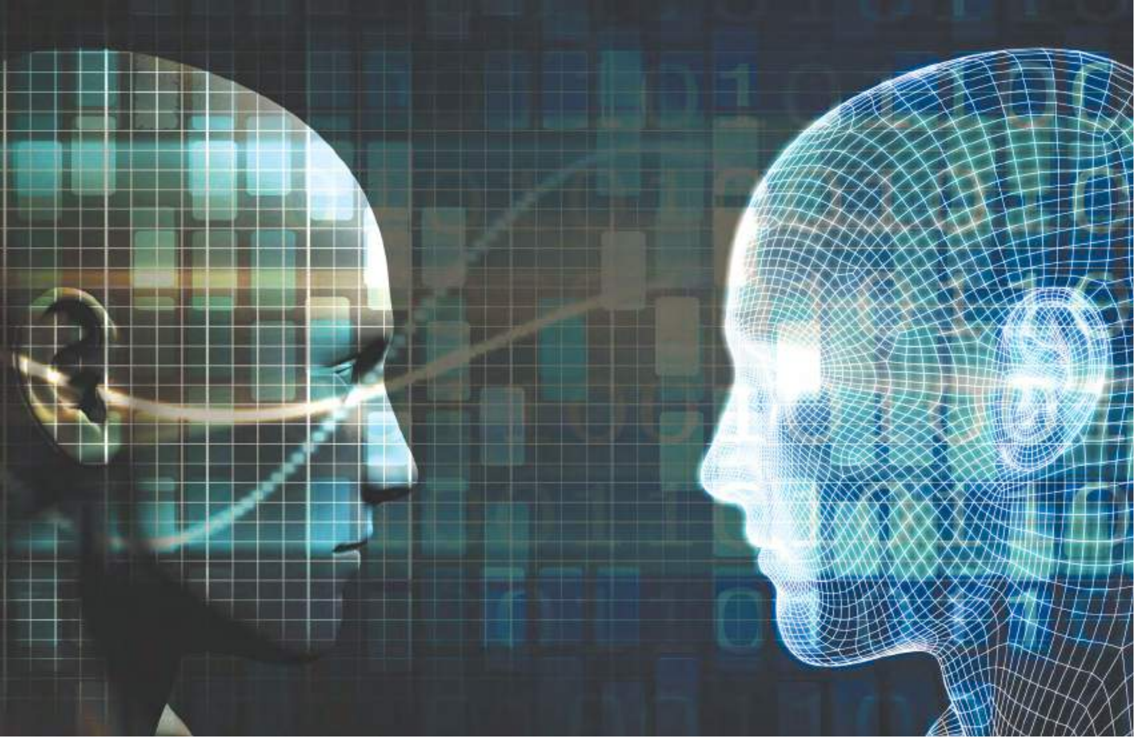

Efficient States

Putting Public-Sector **Productivity** Under the **Microscope**

Juan Gustavo Corvalán
University of Buenos Aires

AN INTEGRATED, UNIFIED GOVERNANCE MODEL COULD IMPROVE THE EFFICIENCY OF THE PUBLIC SECTOR BY THROUGH THE USE OF ARTIFICIAL INTELLIGENCE, WHICH REDUCES BUREAUCRATIC TRANSACTION COSTS. OPTIMIZING STATE PROCESSES IS IN LINE WITH OECD RECOMMENDATIONS. ARGENTINA'S LEGAL SECTOR HAS SHOWN THAT IT IS POSSIBLE TO REDUCE THE TIME SPENT ON RED TAPE AND PAPERWORK BY MORE THAN 75%.

Artificial intelligence (AI) is the most disruptive information and communications technology (ICT) in human history. Intelligent algorithms can now perform as well or better than human beings at a growing number of activities that only our brains were once capable of carrying out. This change is essentially due to radical increases in three interrelated factors: 1) data storage capacity; 2) the speed of processing data and information (big data); and 3) the progressive development of multiple AI systems that recognize patterns to solve problems and achieve objectives. In other words, intelligent algorithms will play an increasingly fundamental role in simplifying environments, optimizing human activities, and maximizing outputs or obtaining other results that would be impossible for us to achieve using our cognitive abilities alone.

In this context, the task ahead is monumental: progress in AI has come at a time of exponential growth[1] which intensifies three characteristic qualities of our era: complexity, uncertainty, and unpredictability[2]. Diagnosing illnesses, taking economic measurements or making economic predictions, designing strategies to promote public policies, protecting employment, and preventing crime are some of the many activities and objectives that fall into these categories. Nobel laureate Daniel Kahneman (2013) calls these "low-validity environments."[3]

AI systems tend to reduce or eliminate the distorted, imprecise judgments and illogical or irrational interpretations that emerge when human brains process data and information. In essence, AI is about handling complexity and uncertainty by reducing cognitive bias and optimizing the handling of the data/information/patterns that underlie human activities and decisions.[4] This optimization is based on reducing times and costs.

However, the end of the 20th century and the first decades of the 21st have been characterized by the transition from a print culture to a digital one. In this process, economic, social, and cultural development has been deeply affected by the use of modern ICTs (computers, the internet, search engines, etc.). We are currently at the dawn of another transition, which has emerged as part of the Fourth Industrial Revolution: we are moving toward new paradigms in which AI will boost human intelligence and give rise to hybrid intelligence. Pattern recognition, the predictions of artificial oracles, the

92%
THE PRECISION RATE OF A CRIME PREDICTION ALGORITHM

radical mutation of our notions of space and time, and the principles of optimization and simplification will transform public strategies and policies that were designed for a world before AI. This is why we need to rethink economic and sustainable development based on hybrid intelligence.

SUSTAINABLE DEVELOPMENT

For a few years now, the UN[5], the OAS[6], the OECD[7], and other international organizations have stressed that ICTs are essential tools for addressing the need for new solutions to the problems of development, economic growth, poverty, and different types of sustainable development.[8] Making ICTs compatible with sustainable development implies tackling four major issues: 1) closing the digital divide; 2) promoting favorable environments; 3) addressing technological readiness; and 4) adopting specific measures to engage with three key ideas: digital inclusion, digital literacy, and inclusive innovation. In recent years, many states have been redefining their approaches based on these principles.[9]

On the one hand, all of these areas can be exponentially increased by using AI systems; on the other, the strategies behind the actions that are being planned for the digital world may eventually be redefined or transformed. This is where we can show, through concrete public organizations—such as the Attorney General's Office of the City of Buenos Aires or the Inter-American Court of Human Rights—how it is possible to drastically change the logic of government processes and procedures that are based on digital systems by transitioning toward ones that combine AI with human intelligence. This would, undoubtedly, represent a qualitative leap in optimizing rights and improving country's development prospects to the maximum. In short, as ECLAC (2018) demonstrates, AI can also be used for economic and social development based on the sustainable development targets adopted by the UN in 2015 (2030 Agenda for Sustainable Development).[10]

THE PUBLIC SECTOR IN THE DIGITAL ERA

With the widespread use and development of the printing press, a print-based form of human identity began to emerge. As a consequence, modern states gradually began to design their organizations based on paper, records, offices, and so on. The space/time duality of a bureaucratic model based on paper and printing is radically different from a digital model based on digital identities. For example, Argentina's National Industrial Registry (RIN) was created in 1972 to evaluate the industrial sector and improve public policy design. The organization includes around just 8,700 companies when there are actually over 110,000 in the country. Furthermore, there are multiple bureaucratic hurdles and complex mechanisms that industries need to move beyond if they wish to access certain government benefits (such as capital goods bonds or the Supplier Development Program, among others). The Argentinian government itself has acknowledged that the process took eight months and was essentially pointless.[11]

Given that this situation is increasing exponentially, the current transition toward digital systems will require other approaches and tools to advance development and optimize citizens' rights. Using one-stop shops or single windows for both online and in-person procedures (which unify citizens' options for accessing government services),[12] promoting simple organizational structures,[13] and good simplification practices are all oriented toward shifting from a single-purpose, decentralized government model to an integrated, unified, all-encompassing one.[14]

The digital public sector is being transformed based on a logic that essentially comes from the private sector: management and organization based on digital platforms that are built around individuals and their habits and preferences, be it for shopping (Amazon and MercadoLibre), transportation (Uber), or in the food sector (Pedidos Ya or GrubHub).

The principles that have emerged through this process were either nonexistent or insignificant within paper-based state bureaucracies. These include: optimizing, updating, simplifying, reducing (costs, complexity, loads, times, etc.), streamlining, robustness, durability, facilitation, flexibility, coordination, harmonization, interoperability, usability, scalability, traceability, and cooperation. Another important factor is making services user-centered so as to make people's lives easier and government services more accessible and inclusive.[15] So far this century, most states have moved toward a digital bureaucracy model based on four main areas of action: 1) switching document and file handling to electronic formats; 2) designing and implementing digital platform-based management systems; 3) online services and online processing of applications and paperwork; and 4) multiple modifications to how administrative functions are organized.[16]

ENHANCED HUMAN INTELLIGENCE

Public-sector data and information tend to be scattered, incomplete, inconsistent, unavailable, or noninteroperable. Indeed, in many cases, such data is not even recorded or stored. In other words, it is not used to add value nor can relevant patterns be obtained from it to allow government administration systems to be optimized and simplified. The example of the RIN that I discussed above is emblematic of how public organizations in Latin America function. In Argentina alone, there are over 3 million public-sector workers (Ministry of Labor, Employment, and Social Security, 2018) who handle millions of documents and files in relation to different economic, social, and cultural rights, among many other things. Sustainable productive development is largely about transforming this lethal combination of data, offices, paper, ink, time, and space.

An intelligent public sector would require a new paradigm: there needs to be a shift from a digital model to one based on hybrid intelligence, which combines human intelligence with AI. This implies a twofold challenge for states as part of a two-part transition.

20
SECONDS:
THE TIME IT TAKES PROMETEA TO HANDLE A COURT CASE

As we move toward an integrated, electronic form of government, we need to rethink the strategies that we use to connect data, information, and information patterns with AI systems and to connect AI systems with human intelligence. The key to this transition toward intelligent states is data governance.

By guaranteeing an interoperable flow of data to which we can apply an AI system, we can radically change crime-prevention policies. Current systems of this type include kernel density estimation (KDE), ProMap, or the PredPol system. In the United Kingdom, a project has been implemented with support from Accenture that uses a predictive AI system to map crime and target use of police resources. The aim of the project is to reduce residential burglaries in urban areas. What is significant about the use of these intelligent algorithms is that they allow as to identify the areas in a city where crime risk is highest. For example, using data from January 2016, it was expected that 248 crimes would be committed in January 2017. When the results were analyzed after the fact, it was found that the AI system had predicted them almost perfectly: 268 crimes had been committed. The artificial oracle was only off by 20 crimes.[17]

This brief example reveals that data flows are the lifeblood of AI. Data needs to be effectively organized and handled and certain information patterns need to be established based on human intelligence so that AI systems can then optimize and simplify their objectives or intended outcomes.

If states start to develop hybrid intelligence models based on AI systems, the public sector's capacity to build on the OECD's recommendations will increase exponentially. These guidelines include promoting ease of access and user-friendly interfaces, facilitating evidence-based decision-making processes, reducing transaction costs, conducting ex-ante and ex-post impact and cost assessments, and identifying/removing/replacing unnecessary, obsolete, insufficient, or inefficient administrative roles or regulations.[18]

AI AT THE SERVICE OF THE STATE

The following example illustrates the potential of this approach. Prometea is an AI that was created in Argentina for use at the Attorney General's Office of the Autonomous City of Buenos Aires. The system was designed and implemented to optimize the legal system, with the aim of exponentially streamlining legal processes to benefit citizens. Promotea is based on supervised machine learning and has proved to be a highly disruptive technological innovation. It was implemented based on the precept that efficient, innovative solutions that impact society as a whole should be government-led initiatives.

This AI system does not learn the law in a human sense but it is capable of things that humans are not, namely reading court cases, predicting outcomes, drafting verdicts, and settling the case in 20 seconds, on average, with an accuracy rate of 96%. In this sense, the system functions similarly to other AI systems, such as Google Translate (although this uses artificial neural networks and Prometea does not). The different intelligent algorithms that Google Translate uses do not actually learn the grammatical structure of different languages (in the human sense of "learning")—instead, they learn patterns that they extract from information and data, which they process at a speed the human brain could never accomplish. Prometea is fairly similar. Although it does not understand laws or jurisprudence, it can nonetheless settle 52% of the less complex cases that reach the Assistant Attorney General's Office for Contentious Administrative and Tax Matters of the Autonomous City of Buenos Aires.

There are three factors that summarize how this unprecedented predictive AI functions and explain why the logic of the system can be transferred to many other bureaucratic government activities (including administrative procedures, queries, permits, licenses, record-keeping, purchases, granting subsidies, dispute settlement mechanisms, etc.)

CASE STUDY: PROMETEA

When the Court of Justice of the City of Buenos Aires (the local high court of justice) sends a file to the Attorney General's Office for a ruling, a human operator accesses the Prometea system and enters the case number. It takes the AI just a few seconds to search for the case number on the High Court of Justice case website. Once the system has found the case, it associates it with another number (in connection with previous action on the case), which it uses to access the Judicial Branch of the Autonomous City of Buenos Aires website (juscaba.jusbaires.gov.ar). On this page, it searches through over 300,000 legal documents and is trained to detect which of these are largely insignificant (for example, instructions to add a certain document to a file or to authorize someone to do something) and which are final rulings by first- and second-instance judges settling different issues.

Once it has located these, the AI system reads them and compares them with more than 1,400 rulings issued between 2016 and 2017 by the Attorney General's Office. After doing all this, within about 15 to 20 seconds, it makes a prediction and provides a written ruling on the case. Some 52% of the issues it handles (with a 96% accuracy rate) are connected with the right to housing, work, and salary-related matters for government employees. The system then asks the user a series of questions which can only be answered by looking at the paper case file (for example, identifying which sheet of paper contains a certain document or statement). This is because the Court of Justice still sends case files in paper form rather than electronically. The whole process is carried out through the AI system, from the first "hello" to Prometea's issuing of the ruling. Once the process is complete, the user can ask Prometea to print the document or download it to their computer.

1. Prometea is fully operational and has settled 96 court cases which, after being reviewed by humans, have been signed and presented before the Court of Justice of the Autonomous City of Buenos Aires. What makes it so disruptive? The main factor is overlapping innovations that, in combination, dramatically cut down on times, costs, and errors. However, it also exponentially increases the benefits for citizens and government employees. According to our calculations, it would take approximately 172 days for human beings to handle 1,000 low-complexity court cases in the legal areas that Prometea has been trained in. Using Prometea, 1,000 cases can be handled in just 42 days without the need to then check for spelling or grammar errors, calculate timeframes and deadlines, or check that certain copies of documents are included in files, depending on regulatory requirements (see table 1).

2. Prometea is also innovating on two other major counts. First, it has been developed based on an integrated screen model which eliminates the need to switch from one window to another to find information access existing documents. The system can be trained to function as a sort of digital bloodhound (for example, you can use a voice command to ask it to search for a law, which it then brings up for you). The usability of the system has been radically optimized through what is known as "intelligence at the interface." Users simply interact with the system by talking to it (Siri-style), or by texting it (as though it were WhatsApp), and the technology solves each problem by connecting to different systems that can respond to users' needs and through learning processes.[19] Prometea is also trained to track timeframes, deadlines, and other basic requirements for legal documents.

TABLE 1
DAYS NEEDED TO HANDLE 1,000 CASES OR FILES

ASSISTANT ATTORNEY GENERAL'S OFFICE FOR CONTENTIOUS ADMINISTRATIVE AND TAX MATTERS OF THE AUTONOMOUS CITY OF BUENOS AIRES	**WITHOUT PROMETEA**	**USING PROMETEA**	**IMPROVEMENT IN EFFICIENCY**
Housing support—Not selfsufficient	160	38	323%
Housing support—Person with a disability	174	45	289%
Housing support—Single person	164	45	263%
Third party subpoena	190	42	357%
1ST INSTANCE ATTORNEY GENERAL'S OFFICE FOR CRIMINAL PROCEEDINGS, OFFICE NO. 12			
Probation (art. 111)	110	26	318%
Expedited trial (art. 111)	145	33	336%
Application for trial proceedings (art. 111)	167	38	338%

Source: Assistant Attorney General's Office for Contentious Administrative and Tax Matters of the Autonomous City of Buenos Aires.
Note: Each working day consists of seven hours. In probation and expedited trial proceedings and applications for trial proceedings, machine learning for decision architecture was used rather than predictive intelligence.

3. There are two other significant aspects to the way Prometea functions. The first relates to predictability, legal certainty, and equality. A key part of generating environments that foster development entails reducing error rates and providing government responses (in this case, legal responses) that are comparable and consistent when circumstances are similar. Both a person demanding the right to housing and a company needing a license to trade want short, uniform response times and fair solutions that do not vary significantly from one case to the next. This implies that standardization is an appropriate place for the state to start designing strategies to reduce both direct and indirect transaction costs.

The second of Prometea's key features is connected to the role of humans. It could be said that the application of this AI system is a form of automation that humanizes automatic processes. In addition to streamlining and speeding up government responses to citizens, it frees people from carrying out routine or mechanical tasks so that they can devote their human intelligence to more complex cases (while also continuing to help or train AI systems to increase their productivity). A significant proportion of the work done at government offices entails copying and pasting texts, numbers, and so on to provide standardized answers or basic solutions that are thought out once and then repeated hundreds or thousands of times by filling in details on forms and complying with requirements for applications for licenses, qualifications, registrations, subsidies, and so on.

For example, certain criminal proceedings (such as drink-driving cases) at the Attorney General's Office require that 39 items be entered or copied 111 times (including age, address, vehicle model, etc.). Using Prometea, each of these items is extracted data is extracted from the system or, in the worst-case scenario, is only entered once and is then automatically replicated by the system throughout the text, depending on the legal document in question. The flow of decisions is designed sequentially so that data entered into one document is automatically transferred to any subsequent documents that are part of the same case. As well as cutting down significantly on processing times, this intelligent design also substantially reduces error numbers.

In other words, one of the paradoxes of public organizations is that many human resources are used for routine mechanical tasks, which often makes them unavailable to work on more complex problems that cannot be handled by AI systems (at least for the time being). Just as computers, the internet, and word processors helped free up time for us to spend on other tasks, weak AI systems will be key to humanizing traditional public services.[20]

In sum, Prometea can be evaluated using the OECD's terms: that is, on the basis of its results and the effects it has on society.[21] This is when the combination of AI and human intelligence really begins to shine. Looking to the future, systems such as Prometea will be key to moving toward a hybrid intelligence paradigm that promotes intelligent governments that foster sustainable development.

NOTES

[1] For more on these issues, see Mumford (1934), Kurzweil (2005; 2012), Bostrom (2014), Mlodinow (2016: 14-17), and Barrat (2013); on the carry-over effect, see Svante (2014).

[2] For more, see Luhmann (2005: 131).

[3] For more, Kahneman (2013).

[4] See Luhmann (1979).

[5] Since the year 2000, the UN has recognized the potential of ICTs in 10 General Assembly Resolutions: Res. A/72/130, July 13, 2017; Res. A/RES/71/212, January 18, 2017; Res. A/RES/70/184, February 4, 2016; Res. A/RES/69/204, January 21, 2015; Res. A/RES/68/198, January 15, 2014; Res. A/RES/67/195, February 5, 2013; Res. A/RES/66/184, February 6, 2012; Res. A/RES/65/141, February 2, 2011; Res. A/RES/64/187, February 9, 2010; Res. A/RES/63/202, January 28, 2009;

[6] See point 1 of resolution AG/RES. 2905 (XLVII-O/17) in OAS (2017) and OAS (2018: 8).

[7] The OECD recommends that governments "employ the opportunities of information technologies and one-stop shops for licenses, permits, and other procedural requirements to make service delivery more streamlined and user focused." See point 5.5 of OECD (2012: 12).

[8] See point 46 of resolution AG/RES. 2880 (XLVI-O/16) in OAS (2016).

[9] For more on these issues in Argentina, see UNDP (2017: 93).

[10] See United Nations (2012: 45) and point 31 of United Nations (2016).

[11] See 27/2018 of the Office of the President of Argentina.

[12] See Communiqué A 6043 from the Central Bank of the Republic of Argentina.

[13] Decree 434/2016, State Modernization Plan, Ministry of Modernization, Argentina.

[14] See United Nations (2012: 85).

[15] See more in Corvalán (2017a).

[16] See more in Corvalán (2017b).

[17] See Gakrelidz (2017).

[18] For more, see OECD (2012).

[19] For more on intelligence at the interface, see Gruber (2008) and http://tomgruber.org/technology/intraspect.htm.

[20] A distinction is made between weak and strong forms of AI, according to how they compare with human cognitive skills. For more, see Kurzweil (2005).

[21] The OECD (2012) recommends, in paragraph 6.4, that "simplification and reform programs should be evaluated for the public value they deliver based on the resources required. Evaluation should focus primarily on the outcomes and affects the society ahead of the quantification of administrative burdens reduced."

REFERENCES

Barrat, J. 2013. *Our Final Invention: Artificial Intelligence and the End of the Human Era.* New York: Thomas Dunne Books.

Bostrom, N. 2014. *Superintelligence: Paths, Dangers, Strategies.* Oxford: Oxford University Press.

Corvalán, J. G. 2017a. "¿Qué hay de nuevo, viejo? Reformas al procedimiento administrativo nacional (Decretos nros. 891/17 y 894/17 (Parte I)." *DPI Cuántico, Diario Administrativo*. 176.

—. 2017b. "Hacia una Administración Pública Digital." *Temas de Derecho Administrativo*. 2: 621–645.

ECLAC. 2018. *Datos, algoritmos y políticas: la redefinición del mundo digital*. Santiago de Chile: UN-ECLAC.

Gakrelidz, N. 2017. "Predicting London Crime Rates Using Machine Learning." *Dataiku,* January 14.

Gruber, T. 2008. "Intelligence at the Interface: Semantic Technology and the Consumer Internet Experience." Paper presented at the Semantic Technologies Conference.

Kahneman, D. 2013. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Kurzweil, R. 2005. *The Singularity Is Near: When Humans Transcend Biology*. New York: Viking Press.

—. 2013. *How to Create a Mind: The Secret of Human Thought Revealed.* New York: Viking Press.

Luhmann, N. 1979. *Trust and Power.* Chichester: Wiley.

—. 2005. *Risk: A Sociological Theory.* New Jersey: Transaction Publishers.

—. 2013. *The Theory of Society*. Stanford University Press.

—. 2018 (forthcoming). *Organization and Decision.* Cambridge: Cambridge University Press.

Ministerio de Trabajo, Empleo y Seguridad Social [Ministry of Labor, Employment, and Social Security of Argentina]. 2018. *Reporte del trabajo registrado.* April 25.

Mlodinow, L. 2016. *Las lagartijas no se hacen preguntas*. Barcelona: Crítica.

Mumford, L. 1934. *Technics and Civilization.* San Diego: Harcourt, Brace and Co.

OECD. 2012. *Recommendation of the Council on Regulatory Policy and Governance.* Paris: OECD Publishing.

OAS 2016. *46th Regular Session of the OAS General Assembly. Santo Domingo, Dominican Republic June 13–15, 2016. Acts and Documents, Volume I*. Washington, DC: OAS

—. 2017. *47th Regular Session of the OAS General Assembly. Cancún, Quintana Roo, Mexico. June 19–21, 2017. Acts and Documents, Volume I*. Washington, DC: General Secretariat of the OAS.

—. 2018. *Un llamado a la acción para proteger a ciudadanos, sector privado y gobierno*. White Paper Series No. 1. Washington, DC: General Secretariat of the OAS and AWS.

Svante P. 2014. *Neanderthal Man: In Search of Lost Genomes*. New York: Basic Books.

United Nations. 2012. *United Nations E-Government Survey 2012. E-Government for the People.* New York: United Nations.

—. 2016. "Multi-stakeholder Forum on Science, Technology and Innovation for the Sustainable Development Goals: Co-Chairs' Summary." Res. E/HLPF/2016/6, June 24.

UNDP. 2017. *Human Development Report 2017. Information for Sustainable Development: Argentina and the 2030 Agenda.* Buenos Aires: UNDP.

DIGITAL INTEGRATION

Latin America is making uneven headway on e-commerce agreements. While countries like Chile, Colombia, Costa Rica, Panama, and Peru are leading the way by actively negotiating bilateral and regional agreements that contain provisions on e-commerce and the digital economy, the Mercosur countries are lagging somewhat behind. However, some progress has been made within the bloc to bring countries' positions together and consolidate digital integration.

45
the factor by which global trade in data has increased since 1980

90
the number of restrictions on data flows in the world

70
the number of WTO countries that are negotiating an agreement on e-commerce

The free trade agreement between Chile and Uruguay, signed in 2016, and the new free trade agreement between Chile and Argentina, signed in late 2017, both of which are waiting parliamentary approval, set a positive precedent for digital integration in the Mercosur. The creation in 2017 of a working group to draft a proposal for a Mercosur Digital Agenda is another positive, albeit overdue, signal given the rise of initiatives such as Caricom's Single ICT Space Integrated Work Plan and the Pacific Alliance's digital agenda.

Number of Free Trade Agreements with Provisions on E-Commerce

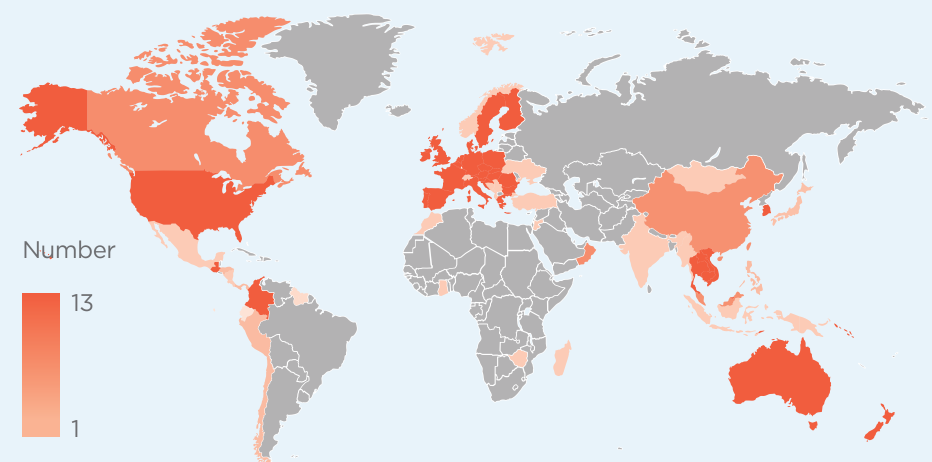

CASE STUDY

When Governments Innovate

Argentina's Policy Design Lab

1,000 QUESTIONS FROM LEGISLATORS HAD TO BE HANDLED IN FIVE DAYS BEFORE THE SYSTEM WAS AUTOMATED

Research in artificial intelligence (AI) is advancing at breakneck speed. Every year, new algorithms break performance records and master new skills. These advances open up opportunities and pose challenges. For governments, the challenge is twofold: they need to stay ahead of the game as both users and regulators of this technology. Two experts from Argentina's Ministry of Modernization, Natalia Sampietro, director of public data, and scientist Martín Elías Costa discussed LabGobAr, the laboratory where algorithms are used to improve public policy design.

What is the aim of the government laboratory at the Ministry of Modernization?

Government innovation labs are dynamic spaces that promote creativity in new public policy designs. These spaces tend to be defined by the fact that their members come from different sectors and tackle problems collaboratively. Unlike traditional government bodies, where changes to processes and policies imply major risks and difficulties, government innovation labs have burst onto the global political scene to take on these risks and foster change. Some of these units test their approaches through experiments, rigorously evaluate the impact of these, and have established themselves as controlled testing spaces for management innovations. Others focus on strengthening the innovation ecosystem throughout the rest of government. Argentina's Undersecretariat of Public Innovation and Open Governance at the Ministry of Modernization has officially established the National Government Laboratory, LabGobAr. The lab is part of Argentina's current modernization and reform process and was started as an interdisciplinary team that brings together different areas of government. It has three core focuses: 1) consulting services, to tackle complex problems using new work processes and methodologies including AI, user-centered design, and other emerging technologies; 2) a design academy, which builds capacities, knowledge, and community for government teams by incorporating 21st-century approaches to management, collaboration, and technology, including training in data science, AI, and lean management; 3) community, to break through the watertight compartments that the state is often divided into and work collaboratively to solve public challenges.

What implications do advances in AI have for government?

The benefits of AI are not just about scale—quality is the other fundamental aspect. These two factors are also true when AI is applied to government. The possibility of applying processes at scale is where AI has the potential to lead to the greatest number of practical outcomes. The examples are endless. Natural language processing can be used to convert unstructured information (texts) into structured information (tables), automatically classify documents, and in contextual search engines. Applications for computer vision include automatically monitoring satellite images for land registries, tax records, and to monitor construction work; automatically detecting traffic accidents using security cameras; checking people's identities using biometric records; and digitizing paper records. In the field of statistical learning, there are predictive models to estimate future demand for the provision of certain state services so as to plan resource allocation appropriately. Other uses include detecting anomalies and frauds, monitoring industry subsidies, and using clustering techniques to find behavior patterns or characteristics that are shared by different groups to design public policies that respond to citizens' individual needs. With regard to the obligation to regulate the use of these new technologies and both the possibilities and risks they bring, the state needs to ensure that the application does not result in discriminatory practices or others that increase inequality.

Could you describe some specific examples of your work in Argentina and the outcomes of this?

There are two recent examples of AI applications that we have been implementing at the Ministry of Modernization's LabGobAr, both of which are open source and are available for the community to reuse and improve. The first is a computer vision project that we are carrying out in partnership with health authorities from the Government of the City of Buenos Aires (GCBA). The government is using egg-laying sensors to monitor the presence of the *Aedes aegypti* mosquito in the city. This entails weekly analysis of photographs of substrates from all over the city to identify whether mosquitos have laid eggs there. The process takes a long time if it is performed by human workers, as they have to manually tag all the photos each week, which creates a bottleneck that restricts the number of sensors that can be installed. We developed a computer vision software package that tags images automatically. The second example is a natural language–processing application that was implemented in partnership with the Office of the Head of Argentina's Cabinet of Ministers. Every month, the team at the Secretariat of Parliamentary and Administrative Relations receives questions that legislators ask in Congress, which it then must classify and forward to the most appropriate person within the public administrative system to be answered. The team has just five days to process around a thousand questions. We intervened by developing a text search engine that uses semantic similarity to tag questions automatically. This tool has improved both response rates and times and has enabled us to identify topics that come up often, even though questions can be formulated in many different ways.

What potential ethical concerns need to be taken into account when designing AI-based public policies?

When implementing predictive or segmentation models that have a direct impact on the form or content of services for citizens, we need to pay careful attention to the possibility of biases being introduced into models; that is, algorithms that inadvertently disadvantage a particular sector of the population. Unlike traditional algorithms, whose operational rules are created and coded by a programmer, AI algorithms infer rules from examples. This is why people often say that an AI model is only as good as the data it has been trained with. If data is a reflection of the realities of life, they will include all the prejudices and discriminatory behavior that exist in society, which will ultimately form part of the final algorithm. Say we want to generate an algorithm that suggests a salary for a job applicant based on their résumé. To achieve this, we would train the model using a large number of real résumés and salaries. The algorithm will attempt to use all the regular features of the data to estimate this and it will notice that gender is an excellent predictor of salary levels, because women tend to be paid 20% less than men for performing jobs with similar tasks and responsibilities. If we base our initial offers on this algorithm, then we would be acting in a biased fashion, widening the gap that we should actually be trying to close. On the one hand, more sophisticated algorithms (such as deep neural networks) are still being actively researched and the way they operate is not entirely clear. They often function like black boxes that receive an input and produce an output without allowing us to understand the underlying process or reasons for this. It would be unacceptable for a state to decide to offer or withdraw a certain service from a citizen without being able to explain why this decision had been reached. We also need to understand that training these algorithms requires huge volumes of data. There is a temptation to draw on personal data, and states possess highly sensitive private information on their citizens. Personal privacy needs to be carefully contemplated before planning any project.

Digital Dividends

Automation poses significant threats to employment. The jobs of the future and how to prepare young people for tomorrow's labor market. The challenges that lie ahead in the fields of education, health-care, and income distribution.

New Jobs

Latin America and Australia

Kate Pounder and Geoffrey Liu
AlphaBeta

AS LATIN AMERICA LOOKS FOR NEW AND MORE STABLE SOURCES OF ECONOMIC GROWTH AND INCOME, ONE POTENTIAL OPPORTUNITY IS INVESTING IN TECHNOLOGIES SUCH AS ARTIFICIAL INTELLIGENCE (AI) AND AUTOMATION. THE INTER-AMERICAN DEVELOPMENT BANK (IDB) HAS IDENTIFIED LOWER LEVELS AND QUALITY OF INVESTMENT AS BEING ONE OF THE MOST SIGNIFICANT FACTORS CURBING LONG-RUN GROWTH ACROSS THE REGION. WITH LATIN AMERICA FORECAST TO HAVE LOWER GROWTH THAN BOTH ADVANCED AND EMERGING ECONOMIES, INCLUDING ASIA, SUB-SAHARAN AFRICA, AND EMERGING EUROPEAN ECONOMIES, STIMULATING NEW AND EFFICIENT INVESTMENT IS A PRIORITY (CAVALLO AND POWELL, 2018).

Investing in AI and automation technologies could form part of the long-term growth solution in Latin America, offering the potential to boost productivity and competitiveness, help existing companies and industries transition, and potentially, create new market opportunities. However, adopting these technologies could also exacerbate regional employment pressures: the McKinsey Global Institute (2018) estimates that in five Latin American countries alone, over 100 million jobs could be impacted by automation. In a region with average unemployment hovering around 8.1% in 2017 and a regional youth unemployment rate of 19.5%, understanding and managing technology-related employment transitions thoughtfully will be critical.

To help frame the choices facing Latin America, this article reviews the potential economic effects of AI and automation, and the policy approaches other countries are pursuing to maximize the economic and social benefits from these technologies. The article begins by taking Australia as a case study and examines whether AI and automation ultimately pose an economic opportunity or a threat to jobs and livelihoods. It then considers how countries around the world are preparing for changes arising from AI and automation technologies, and the lessons and opportunities of these approaches for Latin American countries.

The Australian example shows that the changes brought about by AI and automation can create net economic and social benefits. Greater automation of routine tasks can make work safer, more satisfying, and better paid. New automation and AI technologies can also make workplaces more productive. They can also create domestic and export markets for products and services, and through this, more jobs and higher national income. However, the Australian case also highlights that such changes will not be without cost. Reallocating tasks can displace jobs and increase the number and frequency of job transitions. Automation of routine work is changing the types of skills employees need and increasing the imperative to reskill and upskill throughout careers. Australia has a poor track record of managing previous work and skill transitions, particularly the movement of young people into

80%
OF AUSTRALIANS ARE EMPLOYED IN THE SERVICES SECTOR

full-time work, and the re-employment and reskilling of lower-skilled male workers when traditional jobs disappear. The case identifies that improving the speed and ease by which Australians move between work and skilling is critical to minimize negative impacts from economic and technological shocks.

While the impact of transitions associated with AI and automation is sometimes portrayed in a fatalistic way, Latin American countries and their global counterparts are not powerless in the face of these changes. There are choices governments, businesses, and citizens can make to optimize opportunity and to manage negative shocks. Countries need to understand the different dimensions of change and develop a national strategy to deal with them, to positively influence the change that is coming, particularly for economies such as those in Latin America already undergoing transitions.

TECHNOLOGICAL TRANSITIONS

Australia serves as an instructive case study for the impacts of technological change on an economy. Throughout the 19th and first half of the 20th century, the Australian economy was dependent on mining and resources, agriculture, manufacturing, and construction for growth and employment. However, in the past 70 years, Australia has progressively transitioned to a service-based economy, with around 80% of Australians now employed in the services sector, although sectors such as mining and agriculture still dominate exports. In the last three decades, the economy has also undergone significant structural changes, driven by changes to economic and trade policy, global economic shifts, and technological change. Overall, these shifts have been accompanied by a sustained period of 26 years of economic growth, but in the process, some workers, particularly older, male blue-collar workers have lost jobs and never worked again.

Against this backdrop, there is growing debate in Australia about the rise of a new wave of technologies, epitomized by automation, robots, and AI; the potential social and economic benefits of these technologies; and the risk these technologies may pose to jobs. While Australia is well-positioned to benefit from these changes, the country's recent history is triggering unease about the impacts of AI and has led to questions about how far and fast Australia should push forward with AI-related innovations.

Such debates are not new. Over the centuries, machines have progressively replaced human workers in industries including agriculture, manufacturing, and administration. These technological transitions were also accompanied by political and community concern about the economic impacts of technological change. The Luddite revolts in England during the Industrial Revolution become synonymous with resistance to technological change. In 1895, HG Wells' The Time Machine famously imagined a dystopian and economically divided world where machines had stripped people of the need to work, creating distinct classes of rich and poor.

In practice, previous waves of automation did not cause mass unemployment. Rather, they led to increased prosperity, productivity, and employment, and the creation of new jobs and industries, not the simple displacement of people, as was often feared at the time of the change. This is not to say that technology transitions were painless for those whose jobs were lost, but rather that this pain was concentrated in areas where it was most efficient for labor to be replaced, rather than widespread.

Data from the US also shows that automation is a long-running economic influence, and present rates of job losses from automation are no greater than in the past. Rather, they are comparable to past peaks of technology-driven job losses, and below the peaks seen in the 1970s (see figure 1). The most distinctive feature of the current wave of automation is not that automation is occurring at a higher rate, but rather than it has shifted into service industries.

In the last 25 years, there has been a significant movement of Australian workers from traditional industries to emerging ones. In the process, three quarters of a million laborer and machine operator jobs disappeared, but more than a million jobs were created in care and professional services. Some changes have been driven by the adoption of digital, automation, and ICT technologies. The introduction of ATMs and internet banking, for example, saw 40,000 bank teller jobs shed between the 1990s and 2014. While the transition has not always been painless, the net effect has generally resulted in job creation rather than job loss. In the same period that bank teller jobs were lost, 60,000 financial advisor jobs, an occupation less susceptible to automation because it requires more creative thinking and interpersonal skills, were created (Hajkowicz et al., 2016).

The lesson from these previous changes is that investing in and adopting new technologies is better for growth and jobs in long run, but that in the process the transition will likely create both

FIGURE 1
THE RATE OF AUTOMATION TODAY IS NO HIGHER THAN PREVIOUS PEAKS OVER THE LAST 50 YEARS, BUT THE INDUSTRIES IMPACTED HAVE CHANGED

JOB LOSSES DUE TO PRODUCTIVITY IMPROVEMENT BY SECTOR
% OF EMPLOYMENT LOST EACH YEAR, US DATA

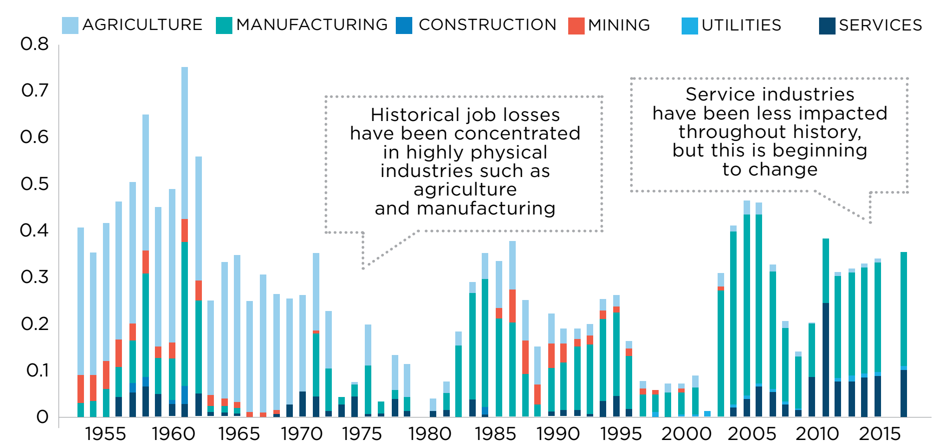

Note: 2011 onwards based on the linear trend for each industry since 1990
Source: Groningen Growth and Development Centre. World-KLEMS database

winners and losers in the job market. How many people fall in the winners and losers categories during the coming AI technology transition will depend on the ease with which the economy creates new jobs and the extent to which workers have the skills and opportunities to transition into them.

AUSTRALIA'S EXPERIENCE

Digital innovations create value in an economy by creating new sources of growth and by increasing productivity. New sources of growth come from both domestic and export income generated by digital industries. Improved productivity arises from investing in digital assets to improve output, and by improving the quality and efficiency of production.

Anticipation of new sources of value from commercializing AI technologies and services is driving increased global AI investment from both the private and public sectors. Excitement about this opportunity stems in part from previous waves of ICT and digital change, which saw the countries and companies that created the technologies and businesses that commercialized them get a disproportionate share of the jobs and profits resulting from them. During these earlier periods, Australia proved to be a better adopter of ICT technologies than a creator and exporter of them. While ICT trade flows have increased by US$5.7 billion from 2000–2017, Australia remained a net importer of ICT through this time (Australian Computer Society, 2017). Australia still lags global technology leaders today. In Australia, technology firms comprise only 3% of Australia's largest 200 companies, and growth in IT market capitalization lags significantly behind the US, which has seen a 168% increase in the technology monthly index since 2011, compared to Australia's index which has risen by 68%.[1]

However, long-term public and private-sector investment in core research capabilities critical to AI and automation, such as robotics, autonomous systems, machine learning, and quantum computing means that Australia starts this wave of technology innovation in a better position to be a technology creator that in previous eras. In 2001, the Australian government created a national ICT research institute, which today aggregates world-class research capabilities in critical areas such as cyberphysical systems, data analytics and optimization, and cybersecurity. In 2000, the ARC Special Research Centre for Quantum Computer Technology was established at the University of NSW and has made a series of critical recent research breakthroughs in the field. The 2016 QS World Rankings for Universities by subject placed four Australian universities in the world's top 50 universities for computer science and information systems (QS World University Rankings by Subject, 2016).

Australia is also in a better position to commercialize the next wave of digital technologies. Australia's economic composition is suited to the kinds of technologies and applications emerging from AI and automation advances. While the last digital revolution through the 2000s focused predominantly on consumer technologies and applications, the AI wave has many industrial uses, which fits more naturally with Australia's domestic economic composition, which is weighted to finance, resources and industrials and energy, and export strengths. In industries such as mining and mining equipment, technologies, and services, Australia has shown already that it can deploy world-leading instances of these. Rio Tinto's Mine of the Future in Western Australia is the largest owner and operator of autonomous haulage systems in the world (Rio Tinto, 2018), and 60% of mining software globally is created by Australian mining services companies (Austrade, 2013).

Australia has a strong starting point and has the potential to benefit from the new value created by AI and automation technologies. The challenge is that Australia is not alone in seeking to build these research capabilities and industries. The value accruing from commercializing new technologies and services have created a global innovation and investment race, and Australia is already falling behind. In particular, the level of business investment in research and development has been stagnant for the last decade, bucking the global trend for national Business Expenditure on Research and Development (BERD) to exceed GDP growth (Innovation and Science Australia, 2017).

AUTOMATION AND GROWTH

New innovations enabled by AI and automation can also deliver productivity benefits to the Australian economy. There is significant scope to increase the gains from automation if Australian firms deepen their investments in productivity-enhancing technologies. If historical trends continue, automation will improve Australia's labor productivity by 8% over the next 15 years. This means automation would drive around one-third of the total expected increase in labor productivity in Australia by 2030. However, Australian firms lag behind global peers in embracing automation. If Australian businesses increased automation investments to match leading countries such as the US, they could add around US$1 trillion to Australia's economic output over the next 15 years (figure 2)

A second source of new value can come from ensuring working hours displaced by machines are reinvested in other tasks or new employment for the minority of displaced workers. If time saved by workers from tasks being automated by machines was deployed to higher-value activities over the next 15 years (rather than simply reducing their work time by 2 hours per week), it could boost Australia's economy by up to US$1.2 trillion in value over that timeframe.

2 HOURS
THE WEEKLY REDUCTION IN TIME SPENT ON ROUTINE TASKS THROUGH AUTOMATION

The impact on jobs as automation and AI replaces labor is central to the equation of whether AI technologies are primarily an opportunity or threat.

Research by AlphaBeta (2017) indicates that most of the impact of automation on jobs will involve augmenting labor, rather than replacing it. The primary impact will change the mix of tasks performed within jobs, rather than the mix or number of jobs in the economy. An analysis of work trends shows that between 2000 and 2015, automation reduced the time the average worker in Australia spends on automatable routine tasks by two hours per week. These include physical tasks, such as lifting or moving items, and cognitive tasks, such as basic information analysis, for example interpreting maps to navigate between destinations or analyzing trends in routine business reporting. If the current pace of automation continues, workers will likely spend another two hours less per week on automatable routine tasks by 2030 (see figure 3).

This suggests that 71% of this expected workplace change will occur within a given job. That is, most of this two-hour change in the weekly work routine will be driven by people doing the same job in a different way, such as shifting from repetitive manual tasks to more complex, interactive ones. The other 29% of the expected workplace change by 2030 will come from workers changing jobs, including where the automation of tasks means fewer people are required to perform a job (AlphaBeta, 2017).

Maximizing the economic benefits of automation is dependent on whether displaced workers are successfully transitioned into new employment. Managing workforce transitions well is critical, as it is costly to individuals and the economy when workers are displaced for long periods of time (i.e. because it increases welfare costs, depreciates skills, and can cause social and economic stress for affected workers). However, as workers will experience different effects from automation, these transitions will take different forms (figure 4).

For many workers, the changes brought about by automation will be positive. These workers (figure 4, top right) will spend more time on problem-solving, creative and strategic thinking, and personal interaction with colleagues, clients, and others. Their work will be more satisfying and better paid, as an hour of nonautomatable work pays 20% higher wages than an hour of automatable work. Jobs will also become safer, as automation frees workers up from routine, manual tasks, which have higher incidents of workplace accidents (AlphaBeta, 2017).

However, older and lower-skilled workers who lose their jobs are likely to be amongst the most vulnerable to negative impacts from transition (figure 4, top left). Australia has a poor track record of helping these workers into new roles when their old jobs disappear. From 1990–2015, nearly one in ten unskilled men who lost their job did not return to the labor force. Today, more than one in

FIGURE 2
AUTOMATION COULD DELIVER A US$2.2 TRILLION DIVIDEND TO AUSTRALIA IF WORKERS ARE TRANSITIONED SUCCESSFULLY AND THE UPTAKE OF AUTOMATION IS ACCELERATED

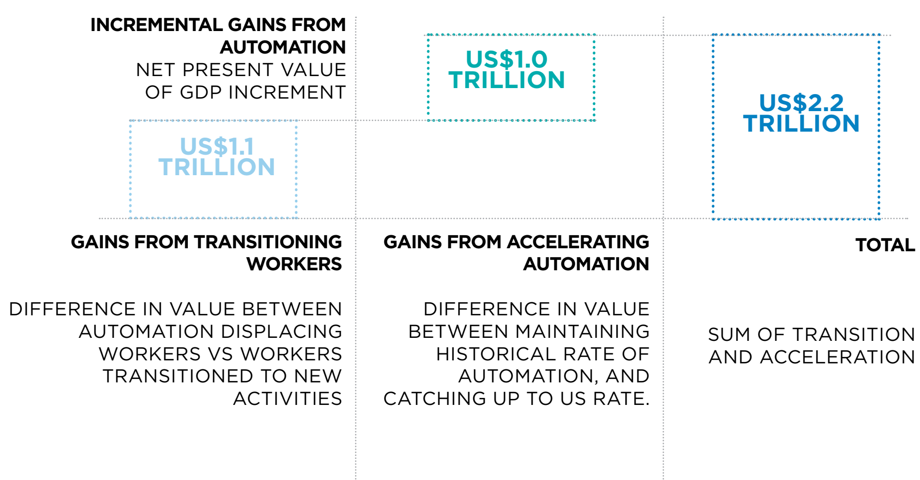

Source: AlphaBeta analysis.

FIGURE 3
AUTOMATION IS CHANGING THE WAY WE WORK, REDUCING THE AMOUNT OF TIME A WORKER WILL SPEND ON ROUTINE TASKS BY UP TO 2 HOURS A WEEK

CHANGE IN TYPES OF TASKS PERFORMED BY AUSTRALIAN WORKERS
AVERAGE SHARE OF TIME SPENT OR WORK ACTIVITY

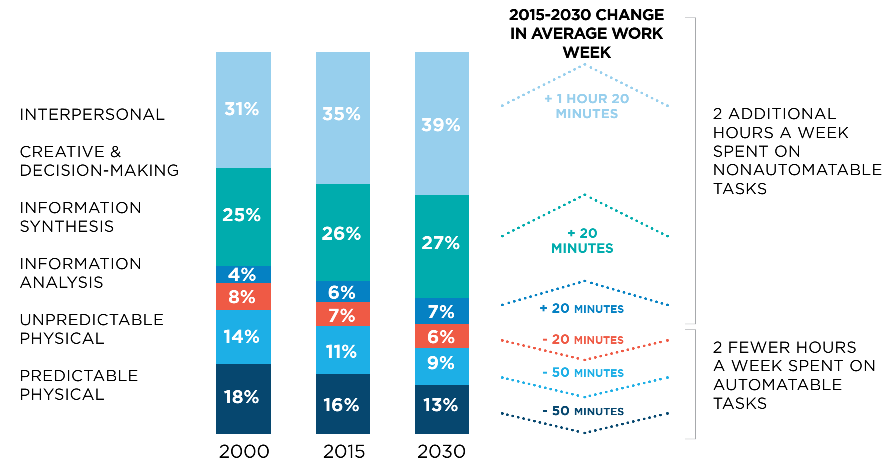

Source: Onet, AlphaBeta Analysis.

four unskilled men still do not participate in the labor market (FYA, 2015).

Future workers, such as students currently in school and tertiary education, need to be equipped for future work and encouraged into training and work pathways associated with growth industries and occupations less exposed to automation. While this sounds straightforward, a significant number of young people are currently training in or entering occupations at risk of automation. Nearly 60% of Australian students (70% in vocational education and training) are currently studying or training for occupations where substantial elements of jobs could be automated by 2030. Around 70% of young Australians are getting their first job in roles that will either look very different or be completely lost in the next 10 to 15 years due to automation (figure 5) (FYA, 2015).

Young Australians face a second challenge when transitioning into full-time work. Working full-time—defined as performing at least 35 hours of paid work per week, either in one or multiple jobs—is an important driver for long-term income, health, and career prospects. However, this is becoming less common for young people. Employment trends show that young people are finding it harder to find full-time work. The share of Australian youths (15–24 years old) in full-time employment has fallen substantially—from 53% in 1980 to around 26% in 2015[2] The proportion of 25-year-olds in full-time employment has eroded from 57% to 51% over the decade that ended in 2016 (Australian Bureau of Statistics, 2006, 2011, 2016). In other words: only every second 25-year-old in Australia is working full-time, and the rest are primarily working part-time, or are unemployed or not in the labor force.

The ease with which workers transition to new work depends on their skills. As the jobs Australians perform are changing, they need a different mix of skills to perform them. The most significant change in skills is the rising importance of enterprise skills and digital and data skills. Training needs will also evolve as more people need to reskill and upskill mid-career to perform the changing tasks in their job and maximize their ability to smoothly manage work transitions. One strategy to help with this is to understand how jobs can be clustered around common skill sets, allowing us to determine the most efficient training pathway for a person to transfer to a new job.

Across jobs and industries, the skills in highest demand are enterprise skills, which include problem-solving, communication, and creative thinking (see figure 6). These skills are essential for the fast-growing “smart creative” jobs in business and professional services, and for the “high touch” interactive jobs in care and household services.

Demand for enterprise skills will increase as the complexity and routine of work increases. By 2030 workers will spend double the time on solving problems and 41% more time on critical thinking and judgment. They will use verbal communication and interpersonal skills 17% more often per week and need to develop a stronger entrepreneurial mindset (FYA, 2017a). Increasing demand for enterprise skills from employers is already evident (see figure 6). Between 2012 and 2015, the number of job advertisements asking for “critical thinking” increased by more than 150%. Demand for workers with “interactive” skills, such as presentation, communication, and teamwork also increased significantly over the three-year period. Enterprise skills are in such high demand that employers are willing to pay an extra US$8,000 for them in early career jobs (FYA, 2017b).

Having enterprise skills also influences the speed with which a young Australian will move into full-time work. Longitudinal analysis shows that courses

FIGURE 4
AUSTRALIA’S POLICY RESPONSE TO AUTOMATION WILL NEED TO BE TAILORED TO DIFFERENT GROUPS OF PEOPLE

GROUPS AFFECTED BY AUTOMATION

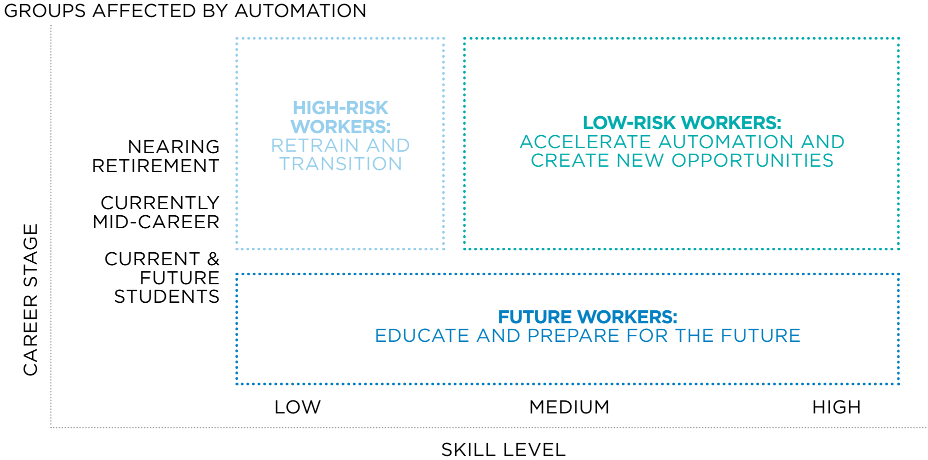

Source: AlphaBeta analysis.

FIGURE 5
70% OF YOUNG PEOPLE CURRENTLY ENTER THE LABOR MARKET IN JOBS THAT WILL BE RADICALLY AFFECTED BY AUTOMATION

BUBBLE SIZE = % OF EMPLOYED YOUTH (15–24 YEARS) IN THAT OCCUPATION

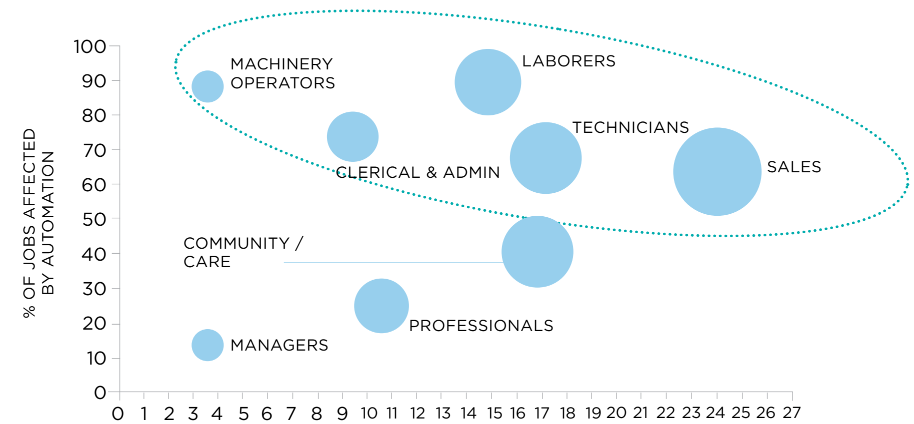

Source: CEDA (2015), ABS, and AlphaBeta analysis.

focused on enterprise skills can substantially shorten the time between leaving full-time education and finding full-time work. A 25-year-old who has completed a course that develops problem-solving, communication, and teamwork enters full-time employment on average 17 months faster than a peer without such training (AlphaBeta, 2018).

A second area of skills growth is digital and data skills. Digital literacy was the highest-growing area of skills demand by employers for early career jobs between 2012 and 2015 (see figure 6). By 2020, 92% of Australian workers will require some form of digital skills in their job. More than half of these workers will need sophisticated skills. Around 8% will need the ability to create technology; critically, 46%—or nearly half the workforce—will require high-level ability to use and configure digital tools and software, and to analyze data (FYA, 2015).

Mid-career retraining is likely to increase as more workers transition between jobs, and more jobs experience a transition in the tasks required to perform them. For the workforce of 2017, 80% of training was delivered before the age of 18. By 2040, only 62% of training will have been delivered before the age of 18 (AlphaBeta, 2018).

Fortunately, many jobs have overlaps in the skill sets they draw on, meaning skills in one job are often “portable” to other jobs. Analysis of Australian job advertisements shows that, on average, people who have learned the skills to work in one job already have the skills to do 13 other jobs (FYA, 2016). These skills can be grouped into clusters and used to inform course design, and career and course advisory services, to help workers switch from one occupation to another.

Australia clearly has many characteristics differentiating it from Latin American countries and economies, and Australian experience cannot be interpolated directly to the Latin American context. However, Australia’s experience may provide some guidance for how Latin American countries can navigate transitions in their own markets. This includes designing responses to responding to economic and technological change in a transitioning economy and lessons for investing in, and managing, the next wave of AI and automation-driven change.

COMPARATIVE STRATEGIES

Like Latin American countries and Australia, many nations are questioning how they should respond to the advent of AI and automation. A growing number of countries have released holistic national strategies for AI and automation, enabling a comparison of different approaches and what this means for emerging regions like Latin America compared with an advanced economy such as Australia.

In the last 18 months, the United States, China, Japan, the United Kingdom, France, and Canada have all released national AI strategies that address economic, social and research aims (see figure 7). These strategies explain how the country will optimize the benefits of these technologies and limit their negative impacts. They set out ambitious plans to build, strengthen, and invest in strategic research areas; to develop ethical and regulatory frameworks that can accommodate difficult questions posed by AI capability; and to minimize negative economic and social impacts, including on employment. While there are many common elements of strategies, each reflects the distinctive economic, social, and cultural context in which they were created, and the unique strengths and opportunities for different countries from AI and automation.

In developing these strategies, governments are acknowledging that they can be purposeful about the values they want AI research activity and technology development to embody, rather than simply react to economic or social shocks.

FIGURE 6
EMPLOYERS ARE ALREADY DEMANDING DIFFERENT SKILLS—THERE IS A RAPID RISE IN DEMAND FOR DIGITAL, CREATIVE PROBLEM-SOLVING, AND INTERACTION SKILLS

GROWTH IN SKILLS REQUIREMENTS LISTED IN AUSTRALIAN JOB ADVERTISEMENTSSTARTFRAGMENT ENDFRAGMENT% GROWTH IN PROPORTION OF EARLY CAREER JOBS REQUIRING THESE SKILLS, 2012-2015

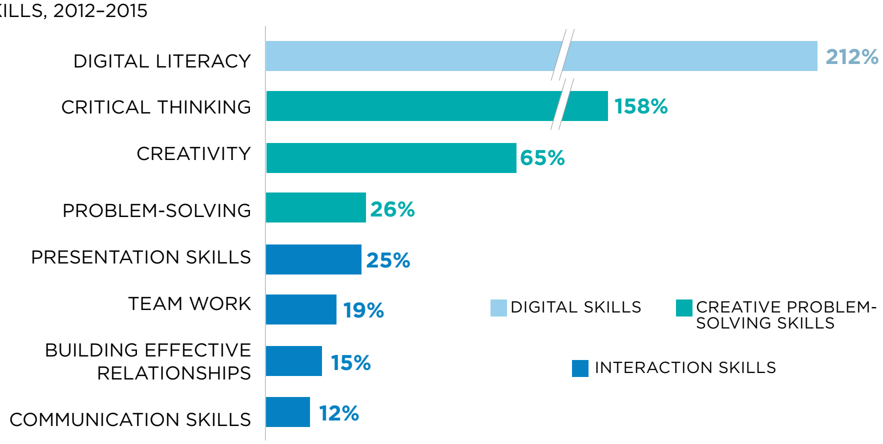

Source: AlphaBeta online job advertisement data. Survey of 4.2 million Australian job ads conducted for FYA (2016).

These approaches highlight that countries have a degree of autonomy over decisions that shape how AI and automation are developed and used in their nation. These decisions shape what countries aspire to achieve from creating or adopting AI and automation technologies, and how they will optimize opportunities and manage risks.

Countries can design an approach that makes sense for their context by analyzing these strategic choices and mapping them to interventions the country can use to optimize opportunities or manage threats.

A comparative review of global AI strategies and approaches globally shows that there are three ways in which countries define AI opportunities and optimize them to their advantage (see figure 8). These are the vision and values the country defines for AI and automation; the economic opportunities and advantages the country can access based on their industry structure and composition, and opportunities to improve social and environmental outcomes.

Australia is already optimizing for opportunity across these dimensions. Its economic composition suits a wave of technological change focused on the business sector, and it has shown it can develop and roll out autonomous systems at scale. World-leading medical research and environmental science capabilities also position Australia well to develop social applications from AI technologies.

The most immediate gap for Australia lies in the lack of an integrated national strategy for AI and automation technologies, defining the goals Australia wants from these technologies, its comparative advantage in realizing them, and its plan to manage negative aspects of economic transition. While some of these areas are covered within other national strategies, they have not been integrated together into a single, coherent vision and strategy. This has prompted both the federal government and state governments to commission research or conduct inquiries to inform Australia's strategic approach to these technologies.

For Latin America, there is an opportunity to be purposeful in thinking about how investment in these technologies could address existing challenges, such as slow long-term growth and lower and less efficient investment, and the areas where Latin American countries have an existing comparative advantage which AI or automation could enhance. There is also the opportunity to determine the broader goals and values important to the Latin American context that would ideally shape approaches to investment and technology adoption.

The experience of other countries also shows how countries can identify enabling opportunities to improve preparedness for AI-led change. A review of national strategies suggests that there are four enabling elements that are key inputs for all countries seeking to be leaders in AI creation and adoption and to avoid the worst shocks from AI-led change. These are people and skills; research and development outputs and access to financing; data assets and infrastructure; and regulatory and ethical frameworks (see figure 9).

Jobs and skills are a focus for Australia. In particular, there is an opportunity to better prepare current and future workers for the evolving labor market. This means using novel data and analytical techniques to get a more immediate and richer understanding of employment trends and create more mid-career training options. Training providers can help students maximize learning and employment outcomes by designing courses that incorporate in-demand enterprise skills that help with faster and more successful entry to full-time work.

In other areas, Australia is already making promising gains. It has strong research and development strengths in critical AI and automation fields, and the Australian government (2018) has recently announced funding to support the development of a national ethics framework to create standards and codes of conduct for developing AI technologies. At the state level, governments are trialing regulatory sandboxes including in areas such as autonomous vehicles.

For Latin America, the question of how to manage transitions in jobs and skills is also important. A useful starting point may be to analyze the existing and future pipeline of workers based on how they may be impacted by AI and automation and segment them accordingly. Comparing the current skills base to projected future employment and skills needs would allow identification of areas where the region has comparative strengths in the talent market and areas where there are gaps to fill. Both pieces of work would enable development of nuanced workforce strategies, and consideration of how the region can tailor training to future shaped by technology, building on the strong increase in educational attainment and skills achieved in recent decades.

Increasing research strengths is another area where Latin American countries could focus. Latin America and the Caribbean increased expenditure on research and development as a share of GDP from 0.56% to 0.77% from 1996 to 2014, but this is still significantly lower than the world average of 2.15% (World Bank, 2018). A prioritized R&D strategy could identify areas of economic opportunity, the areas of AI and automation research and development most relevant to them, and the extent of regional capability in them. This would allow the development of targeted strategies for talent attraction and retention, and investment in research organizations, research infra-

FIGURE 7
COUNTRIES ARE DEVELOPING DIFFERENTIATED AI STRENGTHS AND BRANDS

UNITED KINGDOM
Value proposition: World-leading AI research capability & release of large training datasets

Example: Google's Deep Mind partnered with Oxford University to train AI experts

GERMANY
Value proposition: Advanced manufacturing expertise, skilled workforce, and strong research / industry partnerships

Example: Amazon opened major AI research center in Tuebingen

CHINA
Value proposition: Large and rapidly growing skilled AI workforce, significant investment and vocal government ambition to become global AI leader

Example: China has set the target of creating a US$150bn world-leading AI industry by 2030[1]

CANADA
Value proposition: Strong, aggregated national capability

Example: Funded three national AI research institutes concentrating activity in strategically valuable areas

JAPAN
Reputation for implementing advanced and reliable technologies

Example: Ranked 3rd behind US and China for producing innovative AI research3

UNITED STATES
Value proposition: Large private companies with highly developed AI initiatives

Example: Aside from major firms like Amazon, Google and Apple, the US has over 1,000 AI start-ups[2]

FRANCE
Value proposition: top talent and social and environmental focus

Example: MissionVillani National AI Strategy

SINGAPORE
Value proposition: Government support for smart-city initiatives and experimentation

Example: Smart Nation Singapore initiative

NOTES: [1] Government of China, "A Next Generation Artificial Intelligence Development Plan" (2017); [2] Venture Scanners, "Artificial Intelligence Sector Overview" (January 2018);[3] Times Higher Education (based on Elsvier data), "Which Countries Are Leading in AI Research" (2017), https://www.cifar.ca
Source: AlphaBeta analysis and review.

FIGURE 8
NATIONS CAN SHAPE WHAT AI AND AUTOMATION MEAN FOR THEIR COUNTRY

	KEY CHOICES	OPTIMIZE BY...	MANAGE BY
VISION AND VALUES	What is country's vision for AI and automation in economy and society? What values will define the country's approach to AI and automation? What is the country's unique proposition or brand for AI activity?	Create a national strategy with clear, integrated vision Define national values to shape development and use of AI and automation in country Develop a distinct national AI brand and offering with businesses, government and researchers and actively promote it	Using a national strategy & values to anticipate negative changes, such as employment shocks, and pre-empt or manage them Develop public narratives that address and reduce fears related to change prompted by tech transitionsStartFragment
ECONOMIC OPPORTUNUTY	What are opportunities to improve productivity by adopting AI & automation? How can the country create new sources of growth from AI & automation that build on its comparative advantages? How and where can new jobs and businesses be created? What are the negative economic impacts and risks?	Identify high priority economic opportunities in key industries and target pilot tech trials in them Identify key export and FDI market opportunities and tailor trade and investment attraction programs to realize them Utilize new data sources to create real-time economic indicators for more informed and immediate economic decision-making & evaluation	Review and redesign income, employment and welfare support programs, and regulatory frameworks to keep pace with new work realities
SOCIAL OPPORTUNUiTY	How can governments adopt and accelerate AI and automation adoption? Where can automation and AI improve social outcomes, e.g. in health, public safety, environmental protection? What are potentially negative social impacts?	Develop tech roadmap for government adoption AI and automation Use government procurement to support innovation businesses and applications in priority areas Fund national challenges that concentrate activity & talent on solving biggest social and economic issues	Develop public sector workforce transition strategies Develop AI ethics review capability to identify and manage social impacts

Source: Compiled by the author.

structure, and pilot applications using a mix of temporary and longer-term measures. Given the IDB's findings that lower and less efficient investment has been a brake on long-term growth in Latin America, the strategy could also assess financing gaps for research funding, commercialization of technology, and business investment in new technologies.

INDUSTRIAL CAPACITY AND COMPETITIVENESS

Australia's experience is that it has more to gain than to fear from the adoption of AI. Its world-leading research and industry capability in high-value research fields creating or utilizing new technologies—such as automation, machine learning, robotics, quantum computing, genetics, and medical research—mean it could be globally competitive in commercializing and adopting them. AI and automation also offer new productivity improvement opportunities, which could add up to US$2.2 trillion to the economy. Realizing these benefits is not a given. It will require better preparation of workers for these transitions, including through upskilling. Australia can also increase the coherence of its national vision and strategy for AI, and the coordination of players in the AI ecosystem. However, by starting now, Australia can increase the odds of a smoother transition and greater national economic and social opportunity.

While Latin American countries start from a different position, there are some lessons from the recent experience and approaches of other nations also dealing with the benefits and challenges of technological change which may be helpful to Latin American countries as they chart their own future.

NOTES

[1]AlphaBeta analysis of ASX and NASDAQ.

[2]The terms "young people", "young workers", "young person," and "youth" in this paper refer to 15- to 24-year-olds. For youth employment trends see Australian Bureau of Statistics (1978-2015), Labour Force, cat. no. 6291.0.55.001.

REFERENCES

AlphaBeta. 2017. The Automation Advantage: How Australia Can Seize a $2 Trillion Opportunity from Automation and Create Millions of Safer, More Meaningful and More Valuable Jobs. Sydney: AlphaBeta.

—. 2018. Unpublished analysis

Austrade. 2013. *Mining Software and Specialised Technologies*. Sydney: Commonwealth of Australia.

Australian Bureau of Statistics. 2006. Census Table Builder.

—. 2011. Census Table Builder.

—. 2016. Census Table Builder.

Australian Computer Society. 2017. *Australia's Digital Pulse: Policy Priorities to Fuel Australia's Digital Workforce Boom*. Deloitte Access Economics.

Cavallo, E.A., and Powell, A., editors. 2018. *Latin American and Caribbean Macroeconomic Report 2018: A Mandate to Grow.* IDB.

FYA. 2015. *The New Work Order: Ensuring Young Australians Have Skills and Experience for the Jobs of the Future, Not the Past*. Melbourne: The Foundation for Young Australians.

—. 2016. *The New Work Mindset: 7 New Job Clusters to Help Young People Navigate the New Work Order*. Melbourne: The Foundation for Young Australians.

—. 2017a. *The New Work Smarts: Thriving in the New Work Order*. Melbourne: The Foundation for Young Australians.

—. 2017b. *The New Basics: Big Data Reveals the Skills Young People Need for the New Work Order*. Melbourne: The Foundation for Young Australians.

Government of Australia. 2018. "Budget 2018—New Opportunities and Jobs for Australian Industry." May 8.

Hajkowicz, S.A., Reeson, A., Rudd, L., et al. 2016. *Tomorrow's Digitally Enabled Workforce: Megatrends and Scenarios for Jobs and Employment in Australia over the Coming Twenty Years*. Brisbane: CSIRO.

ILO. 2017. *2017 Labour Overview of Latin America and the Caribbean: Executive Summary*. ILO.

Innovation and Science Australia. 2017. *Australia 2030: Prosperity through Innovation*. Canberra: Australian Government.

Lowe, P. 2017. "Speech: The Labour Market and Monetary Policy." Sydney: Reserve Bank of Australia.

McKinsey Global Institute. 2018. "Where Machines Could Replace Humans—And Where They Can't (Yet)." Tableau Public. Accessed: June 9. https://public.tableau.com/s/gallery/where-machines-could-replace-humans.

QS World University Rankings by Subject. 2016. Accessed: June 9, 2018. https://www.topuniversities.com/university-rankings/university-subject-rankings/2016/computer-science-information-systems.

Rio Tinto. 2018. "Mine of the Future™." Accessed: June 10. http://www.riotinto.com/australia/pilbara/mine-of-the-future-9603.aspx

World Bank. "Research and development expenditure as % of GDP data tool." Accessed: June 10, 2018. https://data.worldbank.org/indicator/GB.XPD.RSDV.GD.ZS.

FIGURE 9
NATIONS CAN SHAPE HOW THEY CAPTURE AI AND AUTOMATION OPPORTUNITIES

	KEY CHOICES	OPTIMIZE BY...	MANAGE BY
PEOPLE & SKILLS	Does the country have the right mix of talent and skills for future jobs? What will job transitions mean for different types of workers? Can education systems deal with changing training needs?	Building enterprise and digital skills development into all tiers of education Better linking work and learning pathways, especially for young people Innovating training to suit mid-career workers' needs	Develop economic and workforce transition strategy based on data about growth areas for jobs and skills Incentivize training in areas of skills scarcity Use skills data-matching to help transition workers to new jobs
R&D AND COMMERCIALIZATION	What core research strengths are needed to be competitive in AI and can the country strengthen or acquire them? How strong are commercialization pathways and collaboration links between research, private and public sectors? Is there sufficient access to financing?	Create a national research strategy identifying priority areas and infrastructure and funding needs Aggregate research in strategic fields through national institutes or centers of excellence Retain top research talent with funding, infrastructure and unique research opportunities	Attract temporary injection talent by hosting global conferences Use skilled migration programs to target top global or expat talent Incentivize applied research and PhD scholarships in strategically important fields
DATA ASSETS & INFRASTRUCTURE	What data assets does the country have and they accessible to the right audiences? Does the country have the right data infrastructure? Is there a data standards and data governance framework?	Make data collected through government initiatives & publicly funded research available Create data standards Invest in data infrastructure	Introduce regulations governing data use, sharing and privacy
REGULATORY & ETHICAL FRAMEWORKS	Can regulatory frameworks deal with AI and automation technologies? What processes does the country have to raise and resolve ethical issues?	Create National Council to proactively identify and advise on AI regulatory and ethical issues Build regulators' access to AI expertiseStartFragment EndFragment Create regulatory sandboxes and test beds	Create mechanisms to engage industry and citizens in ethics decision-making Ensure that the actions and decisions of AI systems are clearly and justly attributable to legal persons

Source: Compiled by the author.

A HISTORY OF COBOTIZATION

1921
The term “robot” is created by the Czech novelist Karel Čapek

1941
Isaac Asimov writes the three aws of robotics:

1) A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2) A robot must obey orders given it by human beings (without breaking rule 1)
3) A robot must protect its own existence (without breaking rules 1 and 2).

1950
Alan Turing publishes his article on thinking machines and the famous Turing Test

1958
IBM produces the first hard drive in history with a capacity of 5MB

1974
Intel produces the 8080 chip

1986
Honda creates the first walking humanoid

1988
The first robot is introduced at Danbury Hospital

1997
DeepBlue defeats world chess champion Gary Kasparov

1998
The first robot capable of interacting emotionally with human beings is created

2005
A driverless car designed at Stanford drives 130 miles on its own

2006
IBM presents Watson, solutions based on artificial intelligence

2011
Apple creates Siri

2015
Amazon creates Alexa, a virtual assistant that connects with household devices

2016
Amazon makes its first delivery using a drone

2017
Google Deep Mind and AlphaGO create an AI that can learn without human input

2018
Uber and Embraer unveil the first flying taxi

Microdata or Macrodata?

An *Alternative Measure for the* Risk of Job Automation

Santiago Chelala
INTAL/IDB

EXISTING ESTIMATIONS OF THE RISK OF THE JOB AUTOMATION DRAW ON MICRODATA TO ESTIMATE PROBABILITIES BASED ON A SUBJECTIVE SELECTION OF THE TASKS THAT ARE MOST LIKELY TO BE AUTOMATED. THIS ARTICLE ANALYSES THE DIFFERENT EXISTING METHODOLOGIES AND SUGGESTS A COMPLEMENTARY MEASURE: A COMPOUND INDEX THAT INCLUDES MACROECONOMIC SERIES IN THE CALCULATION AND ALLOWS PERMANENT MONITORING.

Throughout history, there has never been a shortage of apocalyptic predictions. From the prophecies of Nostradamus to the millennium bug, the end of the world has been forecast on countless occasions for very different reasons. And yet here we still are.[1]

Current pronouncements hailing the end of work are causing panic just as the prophets of old once did. Looming on the horizon is an inescapable threat, one poised to send tremors through the history of humanity: tasks which workers currently earn a wage for performing will be automated using machines that are becoming increasingly efficient and will eventually be cheaper than humans.

Publications on this issue usually try to answer two questions: can a robot perform the work humans currently do? And what will happen if they can? There is no easy answer to either. The first question is hampered by fundamental methodological problems associated with the natural difficulty of making predictions using incomplete information. The second question entails certain aspects of complex socio-economic policy design, such as building new skills and capacities, the regulation of new labor markets, and income distribution mechanisms, among others.

This article focuses on a specific aspect of the first of these two problems: the need for better metrics that would allow us to carry out advance measurements and impact assessments or, as a second-best option, enable us to monitor trends on how technological change impacts our economies.

Technological unemployment is nothing new. In 1930, John Maynard Keynes wrote that "we are being afflicted with a new disease of which some readers may not yet have heard the name, but of which they will hear a great deal in the years to come—namely, technological unemployment."

However, the current pace of innovation and the spread of technology to different aspects of economic life have totally changed the scale on which technological unemployment is unfolding. Brynjolfsson and McAfee (2014) call this new stage "the second machine age." Schwab (2016) prefers a more historical frame of reference, using the term the "Fourth Industrial Revolution," an idea started by Rifkin (1995). Government plans in different countries simply describe the phenomenon as Industry 4.0.[2] These terms all describe the same factors: the Internet of Things, smart cities, big data, driverless cars, artificial intelligence, 3D printing, blockchain, etc. New technologies have become part of the production structure, creating new goods and services and new forms of production, but not necessarily new jobs.

According to the law of transformation, the accumulation of gradual, imperceptible quantitative change necessarily leads to vital qualitative leaps. One radical change that has affected the world of work is population growth. In 1800, there were 1 billion people in the world and it had taken 300 years for the 500 million people that existed in the 16th century to reach this point. In 1920, the global population stood at 2 billion—this time it had only taken 120 years for the population to double again. Today, it does so approximately every 40 years. There are currently around 6 billion people on Earth and this number is slated to increase to 9 billion by 2025. Can our economies generate proportionate numbers of jobs? What sort of jobs will be created and which professions will become obsolete? According to the World Bank (2016), recent decades have seen an increase in the share of occupations that are intensive in cognitive and socio-emotional skills (so-called soft skills), while occupations that are intensive in routine skills have decreased.[3]

Ever since Frey and Osborne (2013) estimated that 47% of jobs in the United States ran the risk of being automated in the following 20 years, a series of studies have proliferated which have attempted to estimate the number of jobs that could be lost, the most vulnerable sectors, the professions that are least likely to disappear, and so on.

Those who use Frey and Osborne's methodology consider that the risk of automation for a given occupation can be transferred from one country to another. Thus, as Arntz, Gregory, and Zierahn (2016) state, the difference in potential job automation rates in different countries is exclusively due to differences in their employment structures. By way of example, if 100% of the jobs in a country were associated with the profession of librarian, which has a 99% risk of disappearing, the risk of job automation in that hypothetical country would be at least 99%. To put it simply, the relevant proportional adjustments are made to account for a different or more complex employment structure.

Using this methodology, Pajarinen and Rouvinen (2014) estimated the risk of job automation in Finland to be 35% while Brzeski and Burk (2015) put the rate at 59% in Germany, to name just a couple of examples. In Latin America, MECON (2016) estimated a risk of 62% for Argentina and Aboal and Zunino (2017) put Uruguay's at 66%. In an extension of the aforementioned study by Frey and Osborne, the World Bank (2016) estimated the risk for other countries, such as China (77%) or Ecuador (68%), using the same methodology.

Far from being immune to criticism, these studies prompted a series of responses that fall into three groups. First, methodological criticism from those who, like Autor (2015), suggest that automation generally impacts specific tasks rather than entire occupations. In other words, a given occupation implies the execution of a diverse number of tasks. In the case of a retail salesperson, for example, these might range from modifying price tags to handling payment or attempting to persuade clients.[4] This approach reduces the risk estimations calculated by Frey and Osborne (2013) and was adopted in Arntz et al. (2016) for OECD countries, who also observed that the same occupation entails different tasks when carried out in a different workplace.[5]

The second wave of criticism zeroed in on the fact that the authors of these studies took a static view rather than a dynamic one. New technologies will also give rise to cobotization (humans and robots working alongside one another in factories) and will come up against regulatory or institutional impediments to automating jobs (such as labor unions). New technologies also give rise to new occupations, as has been the case with data scientists or virtual reality architects. By calculating the risk of automation for a given job without taking into account the creation of new jobs (and the limits on the elimination of existing ones), it may be the case that the negative effect on employment is being blown out of all proportion.[6] Along the same lines, Moretti (2012) argues that each technological job generates a multiplying effect of four new jobs, twice the rate as in traditional industry, due to higher salaries and the propensity of technology firms to form clusters, a factor which is essential to the study of their dynamics in any prospective analysis.

The third source of criticism is historical. Gregory, Salomons, and Zierahn (2016) point out that rather than racing *against* the machines, work races alongside them, in that there is evidence of benefits associated with the increased demand and knowledge spillovers that new jobs generate. Mokyr (2017), meanwhile, observes that the past is a poor guide for predicting the future and that new technologies "will lead to continued improvement in economic welfare, even if these are not always measured in our National Income Accounts." However, measuring this phenomenon precisely poses considerable difficulties. Not even a satellite account for innovation could fully describe probability-related phenomena such as exponential technologies before they are adopted into widespread use.[7]

TABLE 1:
METHODOLOGY AND RESULTS OF SELECTED STUDIES ON THE RISK OF JOB AUTOMATION

	FREY AND OSBORNE (2013)	OECD (2016 AND 2018)	WEF (2016)	WORLD BANK (2016)	MANYIKA *ET AL.* (2017)
METHODOLOGY	Occupation-oriented. O*NET database for the United States	Task-oriented. PIAAC database for OECD countries	Survey of companies from 15 countries in nine economic sectors	Extension of Frey and Osborne into other economies and adjusting for technology adoption lags	Disaggregation of 18 skills used in 2000 work activities in 800 occupations. Estimation of hours used in each activity
RESULTS	47% of jobs at risk of automation in the United States	9% risk of automation on average in 21 OECD countries	5.1 million jobs lost	66% of jobs in developing countries at risk of automation This result is lower if it is adjusted for late adoption of technology	Less than 5% of occupations are fully automatable, but at least 30% of the activities that make up 60% of occupations are technically automatable

Source: Compiled by the author based on Manyika *et al.* (2017).

The so-called productivity paradox described by Roach (1987) shows that investment in innovation and information technology do not move the productivity needle. In other words, even when technology increases, productivity per worker remains the same. Roach's results are a snapshot from just before the internet became widespread, and thus reflect the outlook at the dawn of the IT revolution. Thirty years later, at the dawn of the age of automation, we may be about to witness the rise of a new paradox. This time the issue is undesirable effects not on productivity but on well-being. Automation may translate into lower quality of life for people, greater exclusion, and more unemployment. Why increase productivity if it brings about a more unequal society? How can we better distribute digital dividends to avoid the fragmentation of society in the future? Given this uncertain outlook, any public policy that puts forward alternative courses of action should include the best information available on the evolution of the automation process and the consequences it may have.

A METHODOLOGICAL FRUIT SALAD

The desire to measure a phenomenon as slippery as automation gave rise to a range of very diverse methodologies that brought equally diverse results. Frey and Osborne's (2013) pioneering study argues that there are three bottlenecks to automation, or tasks that cannot yet be automated: those requiring creative intelligence, social intelligence, and perception and manipulation. Frey and Osborne then disaggregate these into nine more specific tasks (such as negotiation, persuasion, originality, and so on) that are used in 702 occupations included in the United States employment database (O*NET). The authors take a subset of 70 jobs and assign them a probability of 1 if they can be automated and 0 if they cannot. The decision in each case was based on consultation with a group of machine learning experts and was thus subjective and ad hoc. The final step in the process entailed generating an algorithm that would predict the automation potential of the 632 remaining occupations included in O*NET, based on the use of the nine task types that make up the bottlenecks.

World Bank (2016) stylizes Frey and Osborne's study to calculate automation risks in different countries depending on their employment structure. That is, the original figures are weighted by the share of different occupations in employment in each country. These results are referred to as being "unadjusted," and the study also includes a calculation that has been adjusted for differences in the rate of technology adoption in poor countries using the technology adoption lag in Comin and Mestieri Ferrer (2013).

The study by Arntz et al. (2016) draws on the database of the Programme for the International Assessment of Adult Competencies (PIAAC) for 21 OECD countries. What makes their analysis different is that it takes a task-based approach that further desegregates the skills used by Frey and Osborne, including microdata on each job relating to tasks such as teamwork or face-to-face interactions. Its results differ substantially from those reached using the earlier approach. While Frey and Osborne (2013) assign a probability of automation of 92% to a retail sales job, Arntz et al. (2016) assign one of just 4%. This difference is due not only to their approach but also to their data source—the PIAAC database allows a much greater level of disaggregation than O*NET.

WEF (2016) is based on a survey of nine industrial sectors in 15 countries and includes 371 firms with a total 13 million employees. At the aggregate level, the analysis shows that new technologies will destroy 5.1 million net jobs.

Using a fairly similar method to Arntz et al. (2016), Manyika et al. (2017), in a study for the McKinsey Global Institute, focus on 18 human skills to estimate the automation potential of 2000 work activities from more than 800 occupations using data from the US Bureau of Labor Statistics. The study was carried out for 46 countries that accounted for 80% of the global workforce. The 18 skills they focus on fall into five categories: sensory perception, cognitive capabilities, natural language processing, social and emotional capabilities, and physical capabilities. An example of their analysis is shown in

The researchers then estimated the level of performance required for each skill used in each of the 2000 work activities on a scale of 1 to 4, where 1 is no performance (or no human skill required); 2 is below median human level; 3 is median human level; and 4 is a high human level of performance. This classification was based on the research group's subjective criteria (as Frey and Osborne's was). The final stage consisted of assigning a given number of hours worked to each activity in each occupation so as to include the risk of automation of the hours actually used for each activity in the probabilistic risk calculation.[8] The result is that less than 5% of occupations are 100% automatable, but at least 30% of the activities that make up at least 60% of occupations have technical automation potential.[9]

The differences in these calculations are noteworthy (see table 1). To cite just just the best-known example, in the United States, Frey and Osborne (2013) and Frey et al. (2016) put the risk of automation at 47%, in contrast with the 10% found by Arntz et al. (2016) and the 14% by Nedelkoska and Quintini (2018), both of which were OECD studies. A simple correlation exercise between the different estimations actually yields negative results (a correlation of -0.35) between the values obtained by Manyika et al. (2017) and Arntz et al. (2016) for the 15 countries included in both samples.

These calculations also include two further problems. On the one hand, they do not allow for periodical comparisons except by recalculating the risk of automation for each activity depending on whatever (nonlinear) technological advances occur. On the other, all the estimations emphasize the impact of technologies on certain occupations/tasks/activities and leave out other relevant factors that form part of the risk of job automation from a broader perspective. These include the population's level of education and the economic structure of a country or its export basket, all of which are relevant when it comes to identifying potential risk factors.

A COMPOUND INDEX OF RISK OF AUTOMATION

There are other automation-related phenomena that are not taken into account in these early studies when they estimate the probability of a given task being computerized. For example, although automation is a risk for all types of tasks, routine tasks are easier to automate and are generally associated with lower levels of education.[10] How does a population's education level af-

FIGURE 1
RISK OF AUTOMATION FOR 37 SELECTED COUNTRIES (IN %)

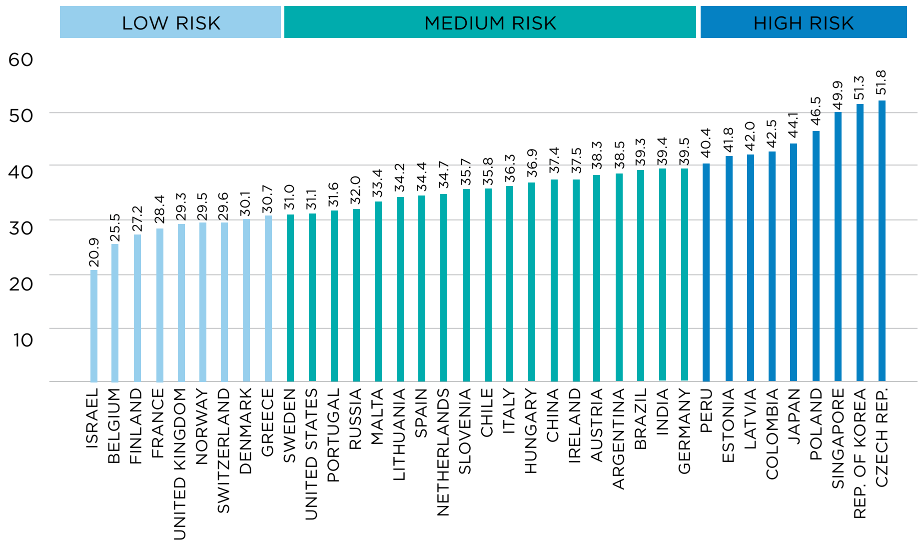

Source: Compiled by the author.

fect the potential for jobs to be lost to new technologies? Another factor that is not contemplated in Frey and Osborne's (2013) calculation is also omitted is the current state of robotization within a given economy, as measured by the number of robots that are already operational within it.

This missing yet relevant information has prompted me to seek an alternative measure for the risk of job automation. This article puts forward a compound index for automation potential which evidently intends to complement rather than replace the estimations analyzed above.

The diverse aspects of automation and the variety of data involved point to the potential usefulness of a compound index that would allow indicators to be aggregated, thus simplifying the analysis and providing economic policy makers with constructive input.

There are a series of well-documented advantages to aggregate indicators. As Jollands, Lermit, and Patterson (2003) argue, "one way to assist policy makers is to develop aggregate indices that summarize the information." Their study looks at a series of aggregate indices that brought solid results in examinations of complex economic phenomena. These include the *Index of Sustainable Economic Welfare*, the *Human Development Index*, the *Unified Global Warming Index*, and many more. Jollands et al. (2003) claim that mathematical simplification is preferable to complexity in these cases and that these indicators are extremely helpful to policy makers because they allow a large amount of information to be summarized succinctly.

Among the potential weaknesses that they discuss, they warn that index aggregation always implies subjective choices and that important information may be lost in the aggregate. Echoing Meadows (1998), they warn: "If too many things are lumped together, their combined message may be indecipherable." The standard criticisms of compound indices run in two directions. First, the choice of the parameters to be aggregated always depends partly on the opinions of the experts

FIGURE 2
VARIATION IN RISK OF AUTOMATION (DIFFERENCE BETWEEN 2017 AND 2014, IN PERCENTAGE POINTS)

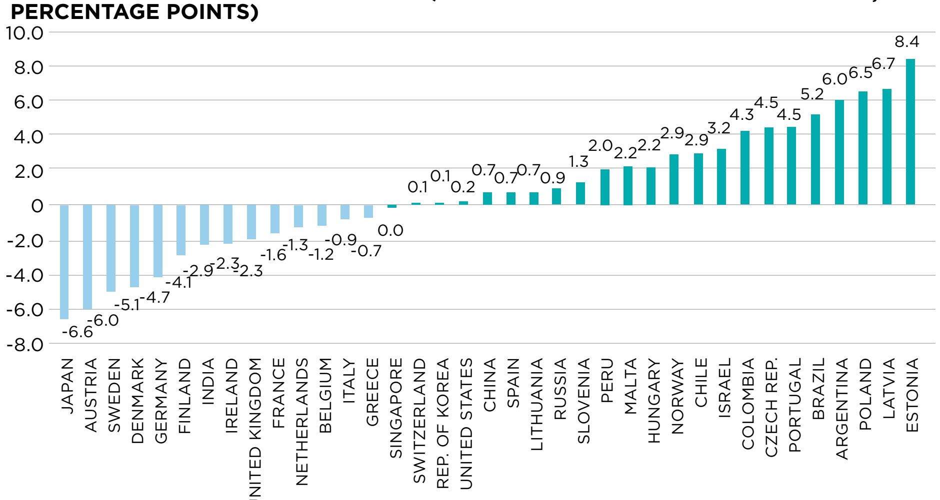

Source: Compiled by the author.

FIGURE 3
COMPONENTS OF RISK OF AUTOMATION FOR THE FIVE COUNTRIES WITH THE HIGHEST LEVELS OF ROBOTS PER WORKER

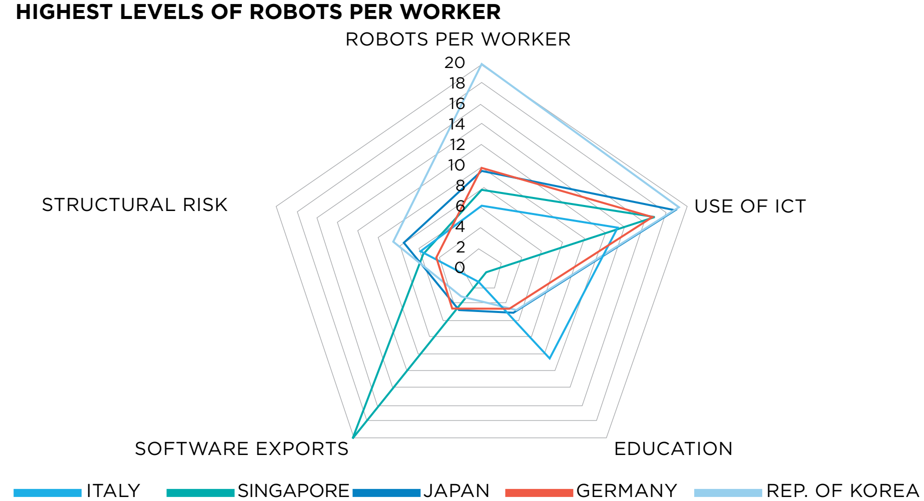

Source: Compiled by the author.

designing the index. Second, it is difficult for aggregate indices to "capture the interrelationships between individual variables." As I observed above, even the most common measures of automation potential are not immune to accusations of subjectivism, while multicollinearity tests can be used to prevent the inclusion of variables that are highly correlated and thus can be thought of as substitutes for one another because they measure the same effect.

The appropriate approach to building compound indices must be based on a clear methodology. As Mazziotta and Pareto (2013) point out, "the heated debate within the scientific community, over the years, seems to converge towards the idea that there is not a composite index universally valid for all areas of application, and, therefore, its validity depends on the strategic objectives of the research."

The OECD (2008) provides a complete guide for building compound indices. The strengths of this type of indicator include the fact that it enables researchers to summarize a set of indices while preserving most variations from the initially released values. The guide thus warns that a prior standardization stage is necessary.

In this case, I have opted for Min-Max normalization, which allows the results of different indicators to have an identical range [0, 1], which coincides with the scale that tends to be used for risk of automation, which has a range of [0, 100]. The normalization criterion is as follows:

$$(1) \qquad I^t_{qc} = \frac{x^t_{qc} - min_c(x^{t0}_q)}{max_c(x^{t0}_q) - min_c(x^{t0}_q)}$$

where x^t_{qc} is the original value of an indicator and I^t_{qc}, is its replacement value after the Min-Max normalization of each series. The different aggregate variables that will make up the compound index thus fulfill the criterion of being scale-invariant, so a unit-based standardization can be carried out and the new values remain within the desired range.

Once the method of standardization has been selected, an aggregation methodology needs to be chosen. The OECD (2008) argues that "by far the most widespread linear aggregation is the summation of weighted and normalized individual indicators," an equation given by:

$$(2) \qquad C_c = \sum\nolimits^Q_{q=1} w_q x_{qc}$$

where $\sum w_Q = 1$, in other words w_Q represents the weight of each variable in the indicator such that $0 \le w_Q \le 1$ for each $q=1, \ldots, Q$, and $c=1, \ldots, M$. This article discusses a compound index in which all the components have the same weight, and leaves an analysis of the results so as to place more weight on some components than others for future studies to consider.[11]

MICRODATA OR MACRODATA?

The index was built for a set of 37 countries (in North America, Latin America, Europe, and Asia) for a four-year period (2013–2016). The index frequency is annual, due to the types of data that technology- and innovation-related variables tend to include.

The variables chosen were connected to automation from a macroeconomic or sector-specific point of view. The aggregation of variables into a compound index thus generates a measure of comparison between the different countries.

Five components were selected to build the index, based on the usual criteria of credibility, coherence, relevance, accessibility, the research group's experience, diversity of aspects observed, and so on.[12] The following variables were included:

1. *Robot stock per worker.* This refers to each country's stock of indus-

FIGURE 4
DYNAMICS OF THE RISK OF AUTOMATION, SELECTED COUNTRIES (%)

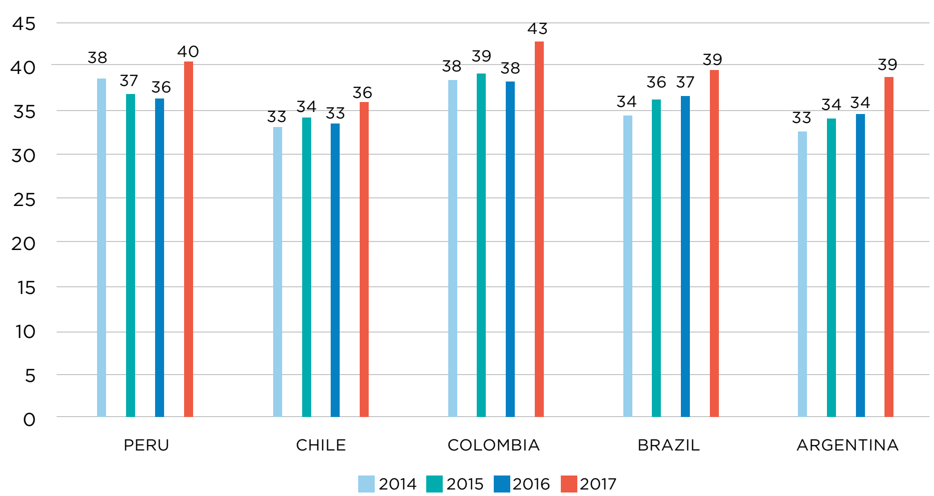

Source: Compiled by the author.

FIGURE 5
DECOMPOSITION OF RISK OF AUTOMATION FOR SELECTED COUNTRIES (%)

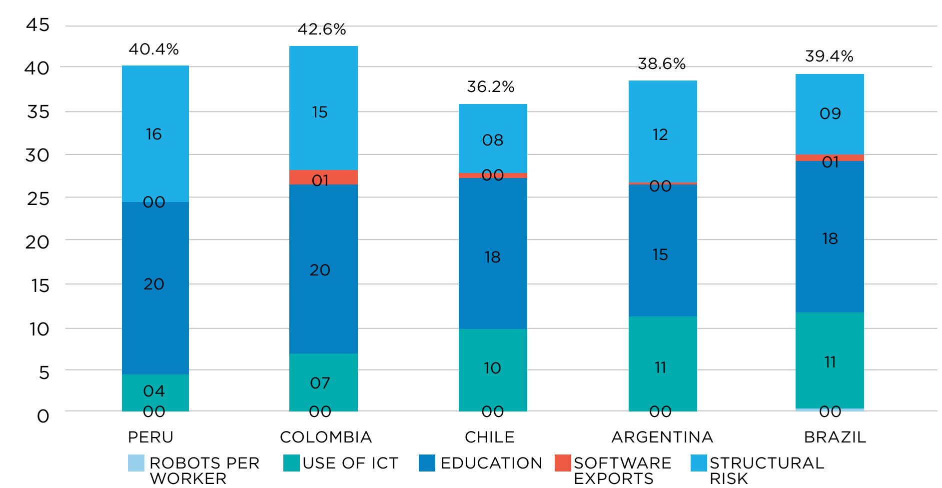

Source: Compiled by the author.

trial robots over time. The assumption is that increased robot density will have a positive effect on the risk of job automation. The sources for this indicator are publications from the International Federation of Robotics (IFR) and the World Bank.

2. *Use of ICTs.* This is an indicator that captures the intensity and use of ICTs. It assumes that greater use of ICTs has a positive effect on the risk of job automation via greater availability of digital automation technology. The source for this indicator is the International Telecommunication Union (ITU).

3. *Education level.* This is an aggregate of variables that include, for example, numbers of science and technology graduates, numbers of students enrolled in higher education programs, numbers of researchers, and education expenditure per country. The source for this indicator was the education section of the Global Innovation Index and the assumption is that the higher the education level, the lower the risk of job automation.[13]

4. *Share of software exports.* This indicator is the share of software exports in each country's total exports, as captured by codes 8523 and 8524 of the Harmonized System (HS). It is assumed that the export baskets of countries whose economies are totally automated will contain high levels of software content.

5. *Structural risk.* This is the weight of employment in sectors that are more susceptible to being automated (where there are more robots per worker). These include agriculture, manufacturing, commerce and transportation, and the hospitality industry, in relation to total employment. These sectors are indicated to be the ones with the greatest risk of automation (Manyika et al., 2017).[14] The greater the weight of the sectors with automation potential, the greater the risk of automation across the economy.

FIGURE 6
DECOMPOSITION OF RISK OF AUTOMATION FOR LATIN AMERICAN COUNTRIES (% OF TOTAL)

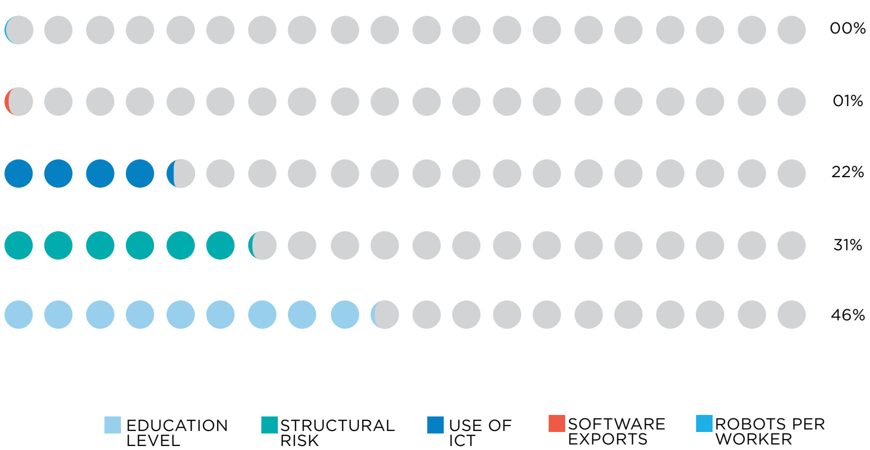

Source: Compiled by the author.

Following a correlation analysis to discard any possible multicollinearity, these five components were weighted identically when the index was calculated so as not to bias the final results toward any of the areas covered. It would be perfectly feasible to change the weighting to place more importance on the present (robot stock) than the future (education) or vice versa. It is important to note that the result will not reflect the absolute risk of job automation but rather the relative risk, given that by normalizing the index components using a range of [0, 100], what is being examined in each case is the relative difference between the countries in question.

RESULTS

Figure 2 shows the results for the 37 countries included in the compound index for risk of automation. At one extreme lies Israel, with the lowest risk (20.9%), which is particularly due to its high levels of education and low structural risk. At the other extreme is the Czech Republic, with the greatest risk (51.9%), which is due to high levels of digitization, high software exports, and high levels of robot usage in the production process. The classification of countries into low-, medium-, and high-risk groups was purely subjective (as tends to be the case): here, 31% was the cut-off line between low and medium risk, and 40% was the threshold for high risk, such that there are nine countries at each end of the spectrum and 19 in the central stretch of the curve.

There is a correlation of 0.57 between these results and those obtained by Manyika et al. (2017) and one of 0.44 with the adjusted version in World Bank (2016).

One of the advantages of this approach is that it allows dynamic observations to be made based on annual

FIGURE 7
RISK OF AUTOMATION AND GDP PER CAPITA, SELECTED COUNTRIES

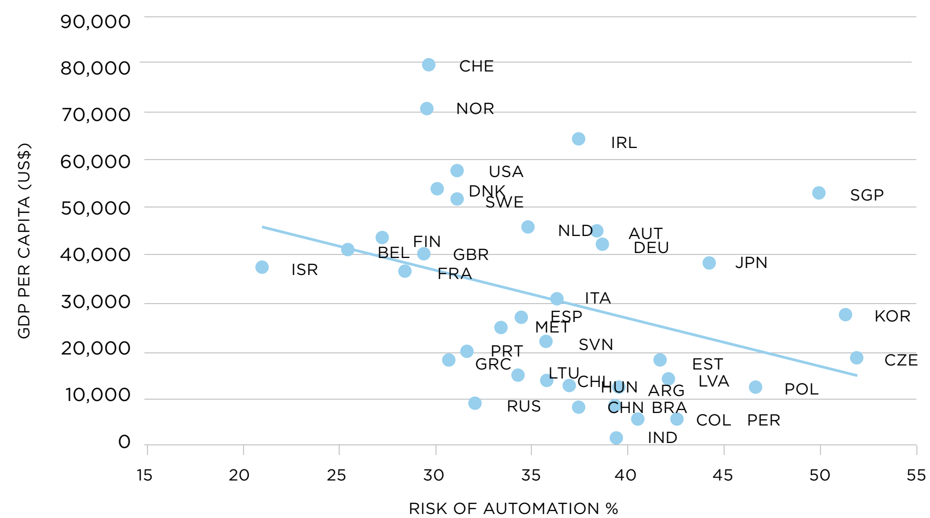

Source: Compiled by the author.

updates of the index components. The percentage difference gained or lost in the last four years are shown in figure 3. Japan, Austria, and Sweden are the countries that have managed to reduce their comparative risk of automation the most, largely by diversifying their productive structure into sectors that are less vulnerable to automation. At the other extreme, Estonia, Latvia, and Poland were the countries that moved up the index most, generally due to a relative decline in the quality of education.

The methodology also allows us to observe the particular composition of risk for each of the countries included. For example, if we look at the five countries with the highest robot stocks per worker, all except Singapore have similar indicator structures, including low levels of structural risk, high levels of education (except Italy), and widespread use of ICTs (figure 4).

As I mentioned above, one of the criticisms leveled at the most widely publicized studies on the risk of automation is that their methodologies prevent the use of prevent dynamic analyses that allow short-term changes in trends to be monitored. The compound index for risk of automation resolves this problem by analyzing the different time series that make it up. With regard to the Latin American countries included in the sample, over the last four years, it can be seen that the risk of automation has increased in Argentina, Brazil, Colombia, and Peru (figure 5).[15]

The decomposition of the risk of automation for Latin American countries shows the most critical factor to be education—the region's levels are relatively low in comparison with the rest of the countries in the sample, and this factor explains 45.8% of total risk. The next-most-significant factor is structural risk, which accounts for 30.6% of the total risk of automation, on average (figures 6 and 7).

The relationship of the compound index can also be compared with traditional economic variables. The following section contains three examples of this: the relationship with GDP per capita, income inequality, and the unemployment rate.[16] First, the compound index shows a negative correlation with GDP per capita of 0.35. Although this information does not constitute an analysis of causality, the empirical evidence shows that employment in countries with higher GDP per capita is at less risk of automation (these also tend to be the countries with the highest education levels; see figure 8).

Furthermore, there is a positive (albeit weak) correlation with the Gini coefficient, one of 0.16. In other words, the countries with the highest Gini coefficients (the greatest inequality) are also those where risk of job automation is greatest (figure 9).

Meanwhile, the correlation of risk of automation with unemployment is, contrary to what one might expect, negative, with a value of 0.24. In other words, the countries at greatest risk of automation, often due to their high current concentrations of robots per worker, also have low unemployment rates, as is the case in Germany, Singapore, or South Korea, to name just a few examples.

This may be because the productivity increases that are generated by digitization or the automation of production counterbalance loss of employment, as some of the literature predicts.[17]

HARMONIZED METRICS

We need more and better measures to monitor the risk of job automation. The variety of results and methodologies that have been used up to now confirm the potential usefulness of harmonized metrics that would allow different countries and different situations

FIGURE 8
RISK OF AUTOMATION AND INEQUALITY, SELECTED COUNTRIES

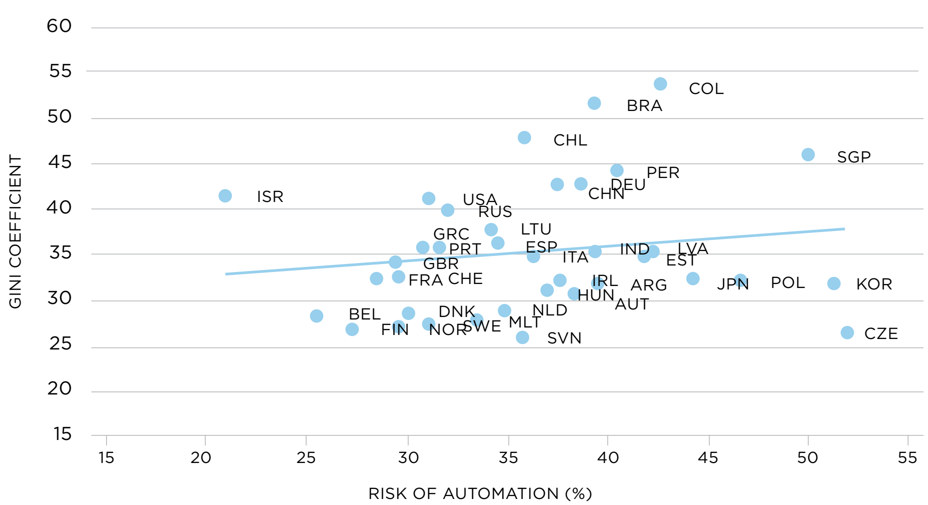

Source: Compiled by the author.

FIGURE 9
RISK OF AUTOMATION AND UNEMPLOYMENT, SELECTED COUNTRIES

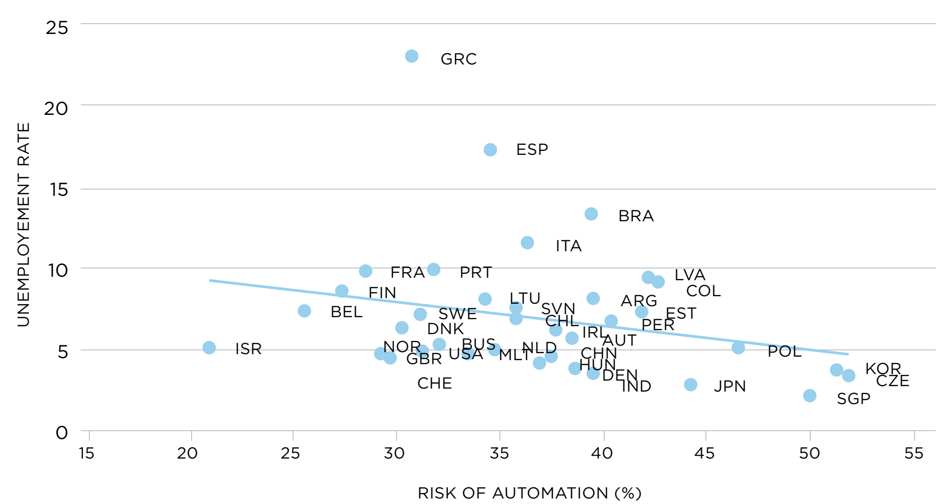

Source: Compiled by the author.

to be compared and monitored over time so as to achieve a reasonable consensus around results.

This article sketches out a possible alternative: a compound index based on other robust indicators. This could also include other components that have not been taken into account in this study, such as data from the private sector on the evolution of employment demand. This is a complementary measure that does not intend to replace microdata-based studies.

The potential advantages of this compound index include its simplicity, the possibility of disaggregating results into the different relevant aspects of automation, and the fact that the calculation can be update periodically as fresh data is released for the indicator components (this could be done on a yearly basis if annual series are used, as is the case in this article).

The results confirm the need for the Latin American countries included in this sample to diversify their exports into sectors that are less at risk of automation, as one third of the potential risk in these countries is currently explained by their productive structures, in which high-risk sectors abound. Alternative sectors these countries could explore include the cultural industries, the orange economy, and knowledge-based services, where the risk of automation remains low.

The reverse empirical correlation found between GDP per capita and risk of automation is a call for developing countries to redouble their efforts to mitigate the negative consequences of the current incorporation of technology into their production processes, as they will be affected more by this factor than developed countries will. Likewise, the negative correlation with unemployment rates raises questions, at the very least, around the bleaker predictions that have been made regarding automation.

NOTES

[1]The author wishes to thank Luca Sartorio and Bianca Pacini for their assistance in organizing the databases to create this index.

[2]Germany, Spain, the United Kingdom, Japan, Canada, and many other countries have launched official strategies for incorporating new technologies into industrial production.

[3]MECON (2016) also highlights this point.

[4]Frank Levy of MIT was far harsher in his methodological criticism when he argued that Frey and Osborne's article "is a set of guesses with lots of padding to increase the appearance of scientific precision. The authors' understanding of computer technology appears to be average for economists (poor for computer scientists)." See http://curriculumredesign.org/wp-content/uploads/Comments-on-Oxford-and-Martin-Study.pdf

[5]Indeed, the skills required by a dressmaker in a Western country would be different to those needed in an Asian country, where more traditional or, in some cases, more sophisticated clothing is the norm.

[6]WEF (2016) attempts to correct this defect by calculating the number of jobs created and lost due to new technologies.

[7]Coremberg and Nofal (2017) look at the need to measure intangible processes. Mokyr (2017) adds that the problem it is complicated by the fact that "the nature of work and the meaning of a job may well change radically as work becomes less and less confined in time and space".

[8]AlphaBeta (2017) contains a similar calculation for Australia.

[9]Manyika et al. (2017) calculate that the adoption of technologies such as personal computers or cell phones took between 5 and 16 years depending on the region.

[10]This claim does not intend to ignore the existence of a vast literature describing polarization and the hollowing out effect, whereby medium-skilled jobs give way to low- or high-skilled jobs (McIntosh, 2013). The situation is similar for white-collar workers who perform skilled jobs that nonetheless run the risk of being automated (accountants, librarians, travel agents, etc.).

[11]An approach based on principal components analysis (PCA) would achieve this. For more, see Jollands et al. (2003).

[12]A complete list of these criteria can be found in OECD (2008, 49).

[13]This assumption is corroborated in the study by Aboal and Zunino (2017), who observed that risk of automation decreases as education level increases.

[14]Sectors where the risk of automation is low or medium were left out of the calculation. This is the case for construction, financial intermediation, real estate, public administration, and education, among others.

[15]The compound index for 2017 was based on data from the 2016 time series; the 2016 index was based on data from 2015; and so on.

[16]The data source used for all three of these factors—GDP per capita, the Gini coefficient, and the unemployment rate—is the latest information available on each country from the World Bank.

[17]This is true of Moretti (2012) and Gregory et al. (2016), among others.

REFERENCES

Aboal, D., and Zunino, G. 2017. "Innovation and Skills in Latin America." *Integration & Trade*. 42: 42–57.

AlphaBeta. 2017. "The Automation Advantage." Working Paper, mimeograph.

Arntz, M., Gregory, T., and Zierahn, U. 2016. "The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis." OECD Social, Employment and Migration Working Papers, No. 189. Paris: OECD Publishing.

Autor, D.H. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives* 29 (3): 3–30.

Brynjolfsson, E., and McAfee, A. 2014. *The Second Machine Age: Work, Progress and Prosperity in a Time of Brilliant Technologies.* New York: W.W. Norton & Company.

Brzeski, C., and Burk, I. 2015. "Die Roboter kommen. Folgen der Automatisierung für den deutschen Arbeitsmarkt" [The robots are coming. Consequences of automation for the German labor market]. ING DiBa Economic Research.

Comin, D.A., and Mestieri Ferrer, M. 2013. "If Technology Has Arrived Everywhere, Why Has Income Diverged?" NBER Working Paper No. 19010. Cambridge, MA: National Bureau of Economic Research.

Coremberg, A., and Nofal, B. 2017. "A Satellite Account for Measuring the New Economy." *Integration & Trade*. 42: 240–254.

Frey, C., Osborne, M., Holmes, C., et al. 2016. "Technology at Work v2.0. The Future Is Not What It Used to Be." Oxford: Oxford Martin School, University of Oxford and Citi GPS.

Frey, C.B., and Osborne, M.A. 2013. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" Working Paper. Oxford: Oxford Martin School, University of Oxford.

Gregory, T., Salomons, A., and Zierahn, U. 2016. "Racing With or Against the Machine? Evidence from Europe." ZEW Discussion Paper No. 16-053. Mannheim: Centre for European Economic Research.

Jollands, N., Lermit, J., and Patterson, M. 2003. "The Usefulness of Aggregate Indicators in Policy Making and Evaluation: A Discussion with Application to Eco-Efficiency Indicators in New Zealand." Working Paper. Canberra: ANU Research Publications, Australia National University.

Keynes, J.M. 1930. "Economic Possibilities for our Grandchildren." In: *Essays in Persuasion.* New York: W.W. Norton & Company.

Manyika, J., Chui, M., Miremadi, M., et al. 2017. "A Future That Works: Automation, Employment and Productivity." McKinsey Global Institute (MGI), January.

Mazziotta, M., and Pareto, A. 2013. "Methods for Constructing Composite Indices: One for All or All for One?" *Rivista Italiana di Economia Demografia e Statistica*. LXVII (2): 67–80.

McIntosh, S. 2013. "Hollowing Out and the Future of the Labour Market." BIS Research Paper No. 134. London: Department of Business, Innovation, and Skills.

Meadows, D. 1998. *Indicators and Information for Sustainable Development*. Hartland, VT: The Sustainability Institute.

MECON. 2016. "Estimaciones preliminares sobre la automatización del empleo en Argentina." *Estudios sobre planificación sectorial y regional*. 1 (1): 4–24.

Mokyr, J. 2017. "The Past and the Future of Innovation: Some Lessons from Economic History." Working Paper. Evanston: Northwestern University.

Moretti, E. 2012. *The New Geography of Jobs*. Boston: Houghton Mifflin Harcourt.

Nedelkoska, L., and Quintini, G. 2018. "Automation, Skills Use and Training." OECD Social, Employment and Migration Working Papers. Paris: OECD Publishing.

OECD. 2008. *Handbook on Constructing Composite Indicators. Methodology and User Guide*. Paris: OECD.

Pajarinen, M., and Rouvinen, P. 2014. "Computerization Threatens One Third of Finnish Employment." ETLA Brief No. 22. Helsinki: The Research Institute of the Finnish Economy.

Rifkin, J. 1995. *The End of Work: Technology, Jobs, and Your Future*. New York: Putnam.

Roach, S. 1987. *America's Technology Dilemma: A Profile of the Information Economy*. New York: Morgan Stanley Special Economic Study.

Schwab, K. 2017. The Fourth Industrial Revolution New York: Crown Business.

WEF. 2016. "The Future of Jobs: Employment, Skills and Workforce Strategy for the Fourth Industrial Revolution." Global Challenge Insight Report, World Economic Forum.

World Bank. 2016. *World Development Report 2016: Digital Dividends*. Washington DC: The World Bank.

Income Distribution

on the Planet of

Superstars

Anton Korinek
Johns Hopkins University

THIS ARTICLE FOCUSES ON THE INTERNATIONAL DIMENSION OF THE INEQUALITY THAT MAY BE GENERATED BY ARTIFICIAL INTELLIGENCE (AI). IF SOME COUNTRIES IN THE WORLD DOMINATE THE DEVELOPMENT OF AI, COUNTRIES THAT ARE LEFT BEHIND WILL SEE CREEPING LOSSES IN TERMS OF TRADE THAT WILL ERODE THEIR STANDARDS OF LIVING. I DISCUSS ALTERNATIVE POLICY OPTIONS TO SOFTEN THE IMPACT ON DEVELOPING COUNTRIES.

If progress in artificial intelligence (AI) continues as predicted, it will usher in an era of unprecedented growth in productivity and output. At present, the greatest scarcity in our economy is the scarcity of labor, as can be seen from the fact that most of the income generated in the economy goes to compensate the labor factor. Even in economies with high unemployment, more than half of all income goes to wage earners. If task after task for which human labor has hitherto been indispensable will be performed by AI, the total amount of output that the economy can produce, and, by extension, total income will rise rapidly.

Some technologists (such as Kurzweil, 2005) even predict a singularity, a point at which machines will become sufficiently intelligent that they master the technology on which they are based and can recursively improve themselves, accelerating the process of technological progress and generating super-exponential growth. The economic implications of such a singularity have recently been formalized by Aghion, Jones, and Jones (2017) and by Korinek and Stiglitz (2017, 2018).

What is there to worry about, you may wonder, if progress in AI ushers in an era of unparalleled abundance? The problem is that every form of technological progress also has a second distinct effect on the economy: aside from increasing the output that can be produced, it leads to changes in relative factor demands and by extension in factor prices across the economy, which generate redistributions between factor owners. In particular, the wages of workers that are substituted by technological progress go down, whereas the incomes of factors that are complements to the innovation—especially the incomes of the innovating entrepreneurs—go up. Figure 1 schematically illustrates these two distinct effects of technological progress on the labor share of output—in the example of labor-saving technological progress on the right-hand side, the income earned by labor declines even though total income goes up significantly.

There is a long history of income redistributions, such as those generated by technological progress, leading to social turmoil. In fact, social turmoil frequently goes hand-in-hand with technological progress. From the very beginning of the Industrial Revolution, when textile production was mechanized in late-18th-century England, technological progress was met by social resistance from the losers—at the time, the artisan weavers who were displaced by machines and who organized the Luddite movement.

Throughout different episodes of history, technological change has exhibited different forms of factor bias. Technological progress in the first three-quarters of the 20th century was overall approximately factor-neutral, implying that the fruits of technological progress were shared quite equally by factor owners throughout the economy, including by workers. However, in the past four decades, the distributive effects of technological advances have been less benign, leading to a significant fall in the relative incomes of workers, as documented, for example, by Karabarbounis and Neiman (2014). In particular, much of the technological progress of recent decades has replaced routine activities that were performed by workers in the middle of the income distribution, leading to growth in both low-paying nonroutine jobs and high-paying cognitive jobs, but hollowing out the middle (see, for example, Autor and Dorn, 2013).

The big question is whether further progress in AI will be a complement or substitute for human labor overall. Although we expect a lot of heterogeneity in the short- to medium-run impact, Korinek and Stiglitz (2017) argue that AI overall is likely to be a labor-saving technology that reduces the relative demand for human labor. In the near future, it is likely that workers along the entire range of the skill distribution will be displaced. Indeed, recent progress in the field implies that machines are getting better both at lower-skilled nonroutine manual activities and in higher-skilled cognitive tasks (such as radiology). There will also be numerous professions in which AI makes workers more productive, especially those who have the skills to create or employ AI tools. For example, programmers who are well-versed in deep learning will continue to experience rising demand. Overall, higher-skilled workers will likely do better in the short-run, since they are typically better able to adapt to changing environments. In the long run, by contrast, if at some point a singularity is reached, machines would be able to replace all human labor.

Another strong reason for why AI is likely to be labor-saving is that it is a digital technology that exhibits the two key properties of any information good: it is nonrival but excludable (we explore this mechanism and its implications in detail in Korinek and Ng, 2018.) Nonrival means that it can be used without being used up—once an AI system is programmed and trained, it can be used by billions at almost zero marginal cost, as exemplified by the online services provided by the Googles and Facebooks of the world. By contrast, a physical good that is used by one person, such as a loaf of bread or a car, cannot be simultaneously be used by billions of others. Excludable means that its owner can prevent others from using it. In the context of information goods, this may be because its owner is keeping the technology secret, or because it is protected by intellectual property rights such as patents.

These two properties imply that AI gives rise to natural monopolies—from an economy-wide perspective, it is most efficient for a single firm to produce the information good and deliver it to the entire economy rather than for multiple competing firms to redouble their efforts and incur the cost of creating the good several times. In the context of AI, this implies, for example, that it is most efficient for a single firm to program the world's search engine, and for a single firm to program the world's social network etc. However, no matter if the resulting market structure is a monopoly or an oligopoly with a handful of players, the side effect is that the resulting firms will have market power. They will charge a markup over their marginal cost of production. This markup serves in part to cover the high fixed cost of the information good, that is, of operating the latest state-of-the-art search engine, but will typically also generate significant monopoly rents in excess of that cost.

As a result of these forces, digital innovation in general, and AI in particular, give rise to an economy of superstars in which a small number of firms or entrepreneurs serve a rising fraction of the market and earn large returns. Every time an AI system displaces jobs that used to be performed by traditional methods using human labor, human wages decline and superstar rents rise, generating greater inequality. Because of this superstar phenomenon, progress in AI has implications for inequality that are much starker than most of the progress that we have witnessed since the beginning of the Industrial Revolution.

One silver lining is that all factors that are complementary to production and that are not easily reproducible or in fixed supply, such as natural resources, will benefit from the economic growth generated by AI. For example, if cutting-edge AI consumes a lot of energy, then energy prices in the economy as a whole will go up, and those who own energy supplies will benefit.

A second silver lining is that even if much of the measured economic income is earned by superstars, consumers may still obtain substantial surplus from free services that are provided by the superstars via cross-subsidization schemes. The business model of many online services, such as search engines or social networks, is to bundle two different services: a free service that is valuable to consumers and that is paid for by what they earn by providing advertising services. Even if all traditional advertisers lose all their business, consumers still obtain surplus from the services that are provided to them for free.

FIGURE 1
EFFECTS OF TECHNOLOGICAL PROGRESS

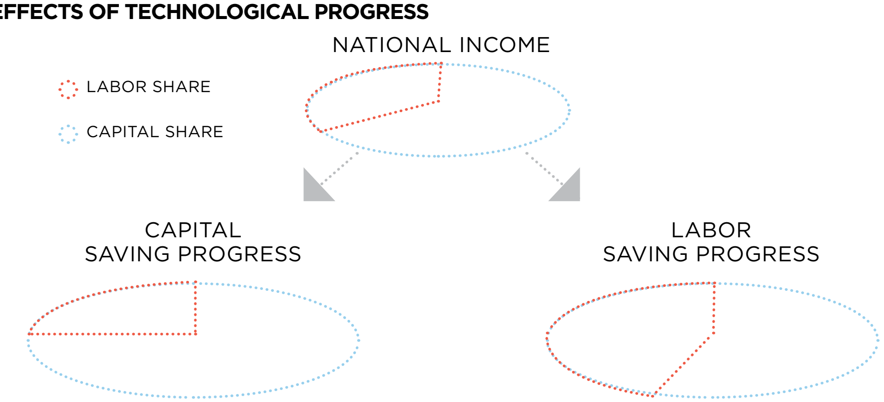

Source: Compiled by the author.

INEQUALITY ACROSS COUNTRIES

So far I have described the forces that may generate inequality as if the redistributions occurred within a given country. However, the potential for greater AI-induced inequality is equally large—if not larger—across countries. The superstar forces that I described earlier imply that innovations in AI may displace labor and hurt workers in all countries around the world, but the resulting gains will accrue only to the superstar in the country where it was developed. This may give rise to a world economy in which there are a few superstar firms in a few superstar countries, with all others lagging behind. Since it is mostly advanced countries plus China and perhaps Russia that are leading the development of AI, this may give rise to a significant amount of reverse redistribution from poor countries to advanced countries, risking a large increase in global poverty.

Technically speaking, the redistribution from poor countries to the superstar countries that dominate the development of AI would occur via deteriorations in their terms of trade. At present, many of the exports of poor countries are intensive in labor, often uneducated or "raw" labor. This is one of the factors that are at greatest risk of being displaced by AI, and if that happens, the relative price of the exports of poor countries would decline compared to their imports, increasing their poverty levels. Furthermore, the creeping terms-of-trade losses that such countries experience may erode their well-being quite silently, since international trade in AI and information goods is not well-captured in national statistics. From the perspective of traditional national statistics, the countries that lose out will, on the one hand, experience slow domestic growth and, on the other, observe that they receive lots of cheap or free digital goods from abroad, but they may not realize that the two phenomena are linked, that is, that they are two sides of the same coin—they experience creeping terms-of-trade losses that silently erode their standards of living.

For theoretical clarity, it may be useful to paint the worst-case scenario: if all labor is replaced by ever-cheaper AI, the factor endowment of a country exporting only labor-intensive goods would be completely devalued by technological change so that the country has, in the limit, nothing useful to export—and hence would not be able to afford any imports. The result would be that the country is effectively in autarky. For most countries, especially small ones that do not have a widely diversified production base, this would be a heavy blow.

As we observed, for countries that export a significant amount of scarce commodities, the terms-of-trade losses in labor-intensive exports may be partially offset by terms-of-trade gains from commodity exports, which may experience price increases as a result of higher demand from superstar countries. However, commodity exports traditionally do not translate into the same number of jobs as manufacturing exports, and natural resources are frequently associated with a "resource curse" that stems from the political difficulty of distributing the wealth generated by resource extraction across society. Furthermore, focusing on primary sector activities such as resource extraction may not position an economy well for future technological progress, in particular future advances in AI.

In summary, the prime question determining the wealth of a country in the age of AI will be to what extent it controls AI technology.

REDISTRIBUTION WITHIN AND ACROSS COUNTRIES

Even if AI will greatly increase inequality, a common response in advanced countries is that the implications may not be too dire since the growth effects of AI imply that the winners can easily compensate the losers. In fact, an economic theorem states that redistribution can ensure that technological progress will always generate a Pareto improvement (that is, it makes everyone better off) within a closed economy (see, for example, Korinek and Stiglitz, 2017).

Furthermore, it is sometimes argued, it would be politically unacceptable for a large part of the workforce to have no income, and so political forces will ensure that workers are compensated if labor is displaced by machines. A policy idea along these lines that is frequently advocated by technologists in Silicon Valley is a universal basic income (UBI) that would pay everybody a certain minimum income, independent of their work status and employability, and that is financed by taxing the winners of technological progress. One of the benefits of such a system, already observed by Milton Friedman, an early proponent, is that it would, at the margin, not distort the decision of its recipients of how much to work since it would be paid no matter how much other income they earn.

However, redistribution is an extremely fraught political issue. In fact, many countries, most notably the US, have seen growing political headwinds to redistribution in recent years. The declining willingness to engage in redistribution at a time when inequality is rising may be partly explained by the growing political influence of superstars.

If redistribution is already difficult within a country, the difficulty of conducting outright redistribution across countries is even greater. For example, few advocates of UBI with any sense of realism mean for the income to be "universal" in the literal sense—they generally suggest a basic income that is distributed universally to all the citizens of their own country. This implies that it will be difficult for countries that experience terms-of-trade deteriorations to receive much compensation for their losses, even if the world economy as a whole grows as a result of technological progress.

However, there is significant scope for policy measures that focus on the root of the problem, that is, the technological forces behind the superstar phenomenon. I observed that digitization creates superstars because it leads to natural monopolies. These natural monopolies are frequently of global scope. If some countries dominate the development of AI, others will increasingly lag behind and may well be worse off, unless policy counters the negative terms-of-trade effects that they experience.

Let me discuss three categories of policy options to deal with the rapid development of AI and the global superstar phenomenon, which are also summarized in table 1.

FOSTERING THE GROWTH OF SUPERSTARS

The prime objective for countries that are participating in the AI race is to generate the global superstar firms that dominate entire industries and constitute global natural monopolies. This allows the countries in question to enjoy the global monopoly rents generated by digital goods which offset any losses from the decline in domestic wages and helps them improve their terms of trade.

One important caveat is that it may not be possible for every country to be a superstar country. There are strong network effects in the AI industry that imply significant headwinds to developing their own AI industry for all countries except the few that are currently leading in the development of this technology. Silicon Valley, for example, offers a large pool of educated workers in the field as well as almost unlimited supplies of venture capital. As a result, competing countries that aim to foster the growth of their own superstar entrepreneurs frequently have to contend with brain drain. Superstar entrepreneurs are highly mobile and have strong incentives to move to hotbeds of digital innovation like Silicon Valley. There is no easy solution to this problem for the countries in question, except to hope that some superstar entrepreneurs will prefer to remain local and that the others will "give back" to their home countries, such as by investing resources in local start-ups or assisting in the education of future superstars.

However, having noted this caveat, let us discuss three policy options that have been proposed. First, education has been the mantra in discussions on how to ensure a country can participate in the knowledge economy in recent decades. However, let me take a more nuanced approach here. Education is certainly desirable for a long list of reasons, including humanistic reasons that are far outside the realm of economics (and which I endorse). Furthermore, if technological progress is skills-biased in the near future, it is desirable to have a better-educated workforce. However, it is not clear that increasing the general level of education in a country is the most effective way of fostering the growth of superstars. Access to a broad talent pool—produced by general education—is certainly helpful, but an important observation is that identifying and fostering the growth of global superstars requires that extra resources are focused on top talent.

A second important factor in the development of superstar firms seems to be a large digital home market, similar to the home market effect in traditional trade theory that Krugman (1980) proposed. This observation suggests significant advantages in the development of AI superstars in countries that have a large customer base at home, such as China and the US. However, small countries can bundle together to form regional digital trading blocs—a goal that places like the EU has taken very seriously and that is advisable for other regions in the world as well.

A third factor is to create a general economic environment that is conducive to the growth of innovative firms. This requires adequate infrastructure, especially cyberinfrastructure, a regulatory environment that avoids stifling growth, and acceptance that risk-taking may sometimes lead to failure.

TABLE 1
POLICY RECOMMENDATIONS

FOSTERING THE GROWTH OF SUPERSTARS	PUBLIC FINANCING OF GLOBAL PUBLIC GOODS	ACCESSING PRIVATELY FINANCED INFORMATION GOODS
• Invest in top talent • Grow the digital home market • Implement policies conducive to innovation	• Public financing of research • Open-access schemes	• Preferential access or compulsory licensing of IP to poor countries

Source: Compiled by the author.

PUBLIC FINANCING OF GLOBAL PUBLIC GOODS

Given that not everybody and not every country can be a superstar, let me return to the basics of economic theory to lay out some policy options from a more general perspective. I characterized the impact of AI-based innovation as creating information goods that are nonrival but excludable.

Many innovations, however, are excludable not by their nature but because we have created intellectual property rights that make them so. Intellectual property rights are not rooted in deep-seated economic laws but are second-best devices for financing innovative activity by conferring temporary monopoly power to innovators. This allows them to earn monopoly rents that constitute an effective way of paying for innovative activity, but a side effect is that it also creates an inefficiency—it leads innovators to charge markups on the goods they sell and thereby sell lower quantities than what would be optimal.

The first-best solution in such an environment would be to publicly finance investment in research and make the resulting innovations freely available to the world. Indeed, a number of powerful technological advances have followed this model—most famously, perhaps, the internet, which was conceived as a decentralized digital network funded by an agency of the US Defense Department.

It is highly desirable to finance as many information goods as possible, including in the field of AI, using public funds and to make them publicly available. Since the benefits of public funding of innovation are global, achieving the socially optimal level of public funding requires some international coordination. A useful analogy is the field of public health, where efforts to coordinate funding for global public goods such as effective drugs or vaccines that are of particular relevance for least developed countries have been highly successful.

However, it is clear that public funding of research and development only goes so far. It is commonly accepted that basic research that delivers very diffuse social benefits is best financed by public funds, whereas innovations that are marketable are often more suited to being developed privately, although this gives rise to the superstar phenomenon that I have discussed.

Interestingly, even a number of information goods that were privately developed have been shared with the public via open-access licenses. For example, several of the superstar firms in the internet sector have made their programming tools for AI publicly available, such as Google's TensorFlow, presumably motivated in part by their desire to expand the number of people familiar with their tools and help them in their recruiting. The more common case, and the one that prevails for most cutting-edge AI applications, is that private firms employ both secrecy and intellectual property laws to make their innovations excludable.

ACCESS TO INFORMATION

For innovations that are better provided by the private sector, offering the information goods produced by superstars at reduced rates or even for free is a good way to compensate countries that lack global superstars for their terms-of-trade losses. There is no deep-seated economic reason why developing countries, for example, should contribute to the financing of innovation in superstar firms that displace their workers at the same rate as countries that host the superstars and reap the majority of the gains. In other words, there is no good reason why intellectual property should enjoy the same rights in countries that host superstars and experience terms-of-trade gains and in those that experience terms-of-trade losses. As I noted earlier, intellectual property rights are second-best mechanisms to finance innovation, and when there are no lump-sum compensations, they should take distributive concerns into account to maximize social welfare.

It is very contentious to advocate preferential access to intellectual property to poor countries that lack superstars, and many recent bilateral trade agreements between advanced and developing countries have in fact steered precisely in the opposite direction. As in all discussions on the topic, the proponents of strong intellectual property rights point to the importance of such rights for the financing of innovation. However, let us be clear: the main reason why the issue is contentious is precisely because weakening intellectual property rights amounts to a large redistribution. This is why armies of lobbyists besiege negotiators for advanced countries to extract concessions on intellectual property rights in trade agreements. To crystallize the issue, the more intellectual property rights a country that lacks superstars assigns to foreign firms, the greater losses in terms of trade and standards of living it will experience.

REFERENCES

Aghion, P., Jones, B.F., and Jones, C.I. 2017. "Artificial Intelligence and Economic Growth." NBER Working Paper No. 23928. Cambridge: NBER.

Autor, D., and Dorn, D. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103 (5): 1553–1597.

Karabarbounis, L., and Neiman, B. 2013. "The Global Decline of the Labor Share." NBER Working Paper No. 19136. Cambridge: NBER.

Korinek, A., and Ng, D.G. 2018. "The Macro-Economics of Digitization and Superstars." Working paper presented at the 5th IMF Statistical Forum. Washington, DC.

Korinek, A., and Stiglitz, J.E. 2017. "Artificial Intelligence and Its Implications for Income Distribution and Unemployment." NBER Working Paper No. 24174. Cambridge: NBER.

—. 2018. "Artificial Intelligence and the Simple Economics of Worker-Replacing Technological Progress." Working paper.

Krugman, P. 1980. "Scale Economies, Product Differentiation, and the Pattern of Trade." *American Economic Review* 70 (5): 950–959.

Kurzweil, R. 2005. *The Singularity Is Near: When Humans Transcend Biology*. New York: Viking.

Expanded *Personalized* Education

Débora Schapira
National University of Tres de Febrero (UNTREF) and
National University of San Martín (UNSAM)

PROGRESS IN EDUCATION-RELATED INTELLIGENCE HAS COME IN THE FORM OF BOTH RESOURCES FOR LEARNING AND EXPERT EDUCATIONAL CONTENT SYSTEMS. THIS ARTICLE ASSESSES THE IMPACT OF CURRENT EXPERIENCES OF USING ARTIFICIAL INTELLIGENCE (AI) IN EDUCATION AND DISCUSSES HOW PUBLIC EDUCATION POLICIES MIGHT BE UPDATED IN THE LIGHT OF THESE NEW TECHNOLOGIES.

Through the use of IT and computer programs, AI provides expert systems that support and facilitate the teaching and learning processes.

Unlike computer-assisted instruction (CAI) models, the main feature of intelligent education systems is that they provide analytical components for undertaking personalized cognitive diagnostics for each student, gathering information on their knowledge level, learning style, psychosociological profile, and motivations. Based on this information, teaching strategies are designed using educational resources that are tailored to each student's characteristics. As these models store multiple forms of data, they provide expert predictive knowledge that can diagnose, assess, and assist the teaching process while also raising the quality of education that each institution provides.

Two types of AI tool are used education: those oriented toward students and those oriented toward teaching staff. The former, intelligent tutoring systems (ITSs), emerged in the 1970s and seek to emulate human teachers. The latter, instructional models, are expert systems that work with teachers to plan and design contents (figure 1).

Professor Jaime Carbonell at the Massachusetts Institute of Technology (MIT) was one of the pioneers of ITSs and developed the SCHOLAR model for teaching South American geography to elementary-school students.

One of the most notable examples of instructional models is the project at the University of Utah[1] led by professor David Merrill, which centers on teaching protocols that guide teachers in designing and organizing knowledge on a certain topic.

Today's ITSs involve new interactions and functionalities, although all have a standard architecture made up of a user interface and three other components: the domain module, the tutoring module, and the student module.

The expert knowledge or domain module contains the essential content on the topic being taught, representing expert knowledge on the issue, providing support and examples, and answering student questions.

The tutor module makes decisions about teaching strategies and automates the teaching process by analyzing differences in understanding between teachers and students and making intelligent decisions based on interaction with the latter.

The student module examines students' different learning styles, problem-solving mechanisms, and knowledge they have acquired previously and makes the necessary inferences for providing the tutor module with feedback based on learning experiences.

The user interface presents the content of each teaching session through text, graphics, and multimedia material. In a nutshell, it allows students to interact with all three ITS modules (figure 2).

ADAPTIVE LEARNING PLATFORMS

There has been remarkable progress in the use of AI in education since the year 2000. The latest innovations are oriented toward using IT functionalities to compile data on the learning process and findings on each student's learning behavior.

These systems and platforms are referred to as "adaptive learning." They seek to detect each student's needs and provide content and activities which they adapt to each individual's learning style. Tracking each student's progress also allows the system to generate personalized assessments. In sum, these models represent more integrated approaches to teaching and learning than traditional ones do.

In the following section, I explore examples of intelligent adaptive learning platforms.

Knewton, one of the largest educational technology companies in the world, began operations in 2008. Various experts refer to it as "the Google of education." It provides digital texts, teaching, online assessment, and learning analytics for each student. Its CEO, José Ferreira, describes it as a "tried and true platform that fosters adaptive learning solutions from the world's major publishers and education companies." It has entered into partnerships with Pearson PLC (UK), McGraw-Hill (US), Cengage Learning (US), Santillana (Spain), Le Livre Scolaire (France), Malmberg (Netherlands), Gakken Educational Co. (Japan), Studentlitteratur (Sweden), UTH Florida University, and Microsoft and Arizona State University for personalized math tutoring.

Since 2009, McGraw-Hill's education department has been developing different intelligent platforms for higher education.

· SmartBook is a textbook that uses adaptive technology for reading and learning, has three million users, and can be used by students and teachers. When texts are presented to students, SmartBook detects their weak areas and text comprehension skills, highlights inconsistencies in their answers, and offers to review the content in question. It helps teachers include innovative educational technology tools in their classrooms, give students guidance on choosing e-texts, and encourages them to start reading.

· LearnSmart is another McGraw-Hill platform that can be used in combination with SmartBook and is available for 90 areas of study. It prepares students for exams based on key topics covered and the behavior of millions of users. Its individualized approach to learning has increased pass rates by 15%.

15%
THE INCREASE IN EXAM PASS RATES USING LEARNSMART

Smart Sparrow, founded in 2010, is an adaptive learning platform that was started at the University of New South Wales, Australia. It is primarily designed for teachers to use when planning and designing their classes. It includes a very user-friendly tutorial that facilitates the creation of content in different screen and graphics formats and also includes other highly attractive components. It allows users to insert questions on the text and includes utilities that simulate lab work, particularly for the medical sciences. It reinforces what students have learned through didactic games that support the teaching process.

Carnegie Learning, founded in 1998, is a leading intelligent research and teaching platform for sixth-, seventh-, and eighth-grade math that also includes textbooks. It was originally developed by researchers at the Robotics Institute at Carnegie Mellon University in Pittsburgh and its products are currently sold by the company of the same name. It was chosen by the US Department of Defense Educational Activity (DoDEA) as a supplementary provider of math software for 24,000 middle- and high-school students on US military bases in the Americas, Europe, and the Pacific.

Geekie, launched in Brazil in 2016, is an adaptive learning platform for elementary schools. The platform is accredited by Brazil's Ministry of Education, is used in 5,000 schools, and has 5 million users. It has been recognized as one of the world's top five innovative education initiatives and is a member of the UNESCO Associated Schools program. Its advantages include preparing students for the National High School Examination (ENEM), which is

FIGURE 1
AI APPLICATIONS IN EDUCATION

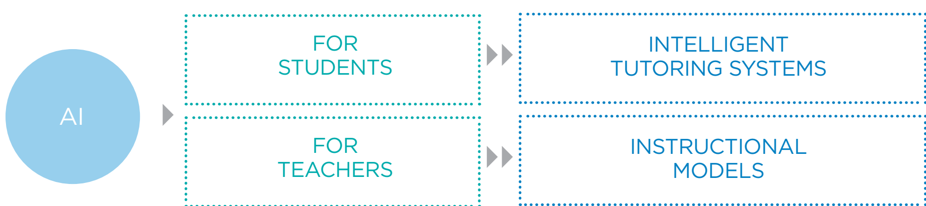

Source: Compiled by the author.

FIGURE 2
INTELLIGENT TUTORING SYSTEMS

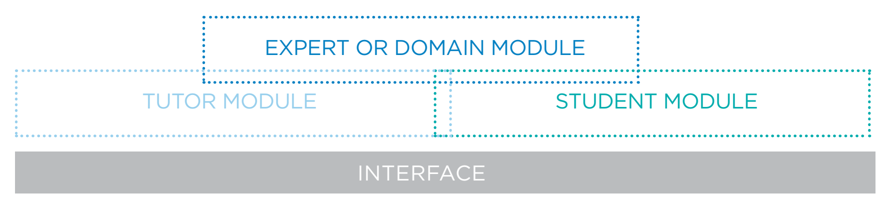

Source: Compiled by the author.

5 YEARS
THE LENGTH OF AN UNDERGRADUATE DEGREE IN ROBOTICS IN ARGENTINA

compulsory for school leavers wishing to enter public higher education establishments in Brazil.

NEW TRENDS

These intelligent platforms' capacity for growth is unimaginable. Their compatibility with traditional e-learning applications marks another new trend. In the following section, I look at some examples of these developments.

Learning Management System

To give classes, e-learning systems currently draw on different learning management systems (LMSs). The best known of these are the Moody and Blackboard platforms. These LMSs are perfectly compatible with adaptive learning systems. Different forms of partnership and collaboration have emerged between companies and educational institutions that each function as links in a value chain, thus closing the gap between the benefits that each system offers.

At the institutional level, several universities in the United States have developed models of this sort. What is interesting about these developments is that they are projects that originated in national education systems, unlike those that come from the education industry.

To illustrate these models, I will focus on those developed at the University of Georgia, Oregon State University, Portland State University, Colorado Technical University, and Open Universities Australia. The latter two have developed the Intellipath and Personalised Adaptive Study Success (PASS) platforms, respectively, for teaching algebra and monitoring student performance.

Adaptive Learning

Pearson PLC runs MyLab and Mastering, an intelligent platform that works with different university courses and that currently has 11 million users. It provides teaching staff with education technology for creating adaptive tutorials, online text, and assessments.

Santillana has developed an adaptive learning product for teaching math at high school level which it has pioneered in Spain and Latin America under the name *A_2O, Aprendizaje Líquido* [L_2O, Liquid Learning]. The experimental version has been used by over 70 algebra teachers and 1,000 students between the ages of 11 and 13 from different countries.

FIGURE 3
ADAPTIVE LEARNING PLATFORMS

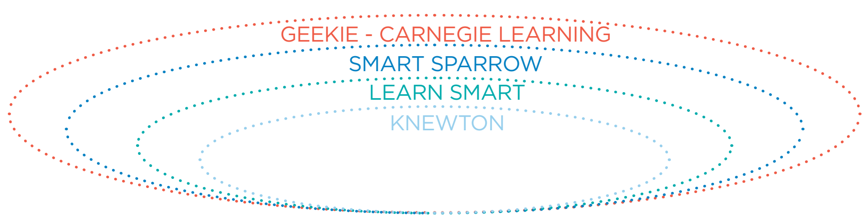

Source: Compiled by the author.

FIGURE 4
TRENDS IN AI

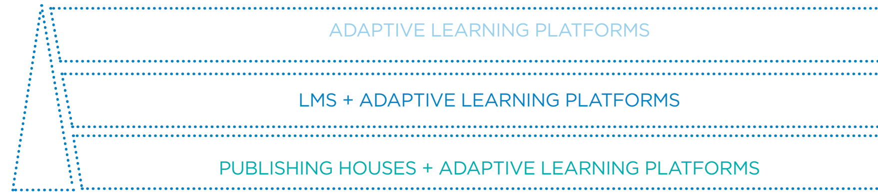

Source: Compiled by the author.

New Education Systems

Another point of interest is the importance of AI and robotics degrees among university programs in Latin America and the rest of the world.

Although it is true that specific degree names may vary, these fields of study are widespread and are indicative of the level of innovation that each country's scientific and education systems are prepared to aim for.

One of the most commonly used names for these programs is "mechatronics," the origin of which alludes to three traditional engineering fields: mechanical engineering, electronic engineering, and computer science (IT). The aim of these courses is to train graduates to design and construct complex products and machinery.

How do AI programs articulate with education as a whole? These programs develop students' skills at designing and programming IT solutions for intelligent systems, particularly for the education system and the adaptive learning subsystem. These applications automate tasks and simplify educational dynamics, and thus provide support for both teachers and students.

It is no coincidence, therefore, that countries at the cutting edge of technological innovation with high education performance standards are including these specialized fields of study in their academic programs.

To illustrate this, I have selected some of the global leaders in this area: Canada, Finland, Israel, Japan, the United Kingdom, Germany, and the United States; and, in Latin America, Mexico, Colombia, Chile, Peru, Argentina, and Uruguay (see figure 5).

IMPACT ASSESSMENT

Educational Initiatives (EI) is a high-profile education organization in India that promotes better learning practices and provides the country's public and private schools with adaptive learning systems. One of these is Mindspark, which aims to help teachers and students and assess the latter's progress on mathematics, language, and science between the third and ninth grades. The demands and effectiveness of this test are comparable with the high-school-level qualification taken by students in Britain, the General Certificate of Secondary Education (GCSE).

Studies in India have shown that the exam results in schools where Mindspark has been implemented as an adaptive learning tool were significantly better than those in schools us-

FIGURE 5
AI DEGREES

COUNTRY	UNIVERSITY	DEGREE PROGRAM	DURATION	LEVEL	PUBLIC	PRIVATE	OBSERVATIONS
URUGUAY	TECHNOLOGICAL UNIVERSITY OF URUGUAY (UTEC)	MECHATRONIC ENGINEERING	6 YEARS	UNDERGRADUATE	X		
CHILE	UNIVERSITY OF CHILE	CIVIL MECHANICAL ENGINEERING	6 YEARS	UNDERGRADUATE		X	
	FEDERICO SANTA MARÍA TECHNICAL UNIVERSITY (USM	UNIVERSITY-LEVEL TECHNICAL QUALIFICATION IN ROBOTICS AND MECHATRONICS	2.5 YEARS	TERTIARY EDUCATION		X	
	ANDRÉS BELLO UNIVERSITY (UNAB)	AUTOMATION AND ROBOTICS ENGINEERING	4 YEARS	UNDERGRADUATE		X	
COLOMBIA	JORGE TADEO LOZANO UNIVERSITY, BOGOTÁ (UTADEO)	INDUSTRIAL APPLICATIONS OF AUTOMATIC PROCESS CONTROL	N/D	UNDER-GRADUATE		X	DIPLOMA
	PONTIFICAL BOLIVARIAN UNIVERSITY (UPB)	SPECIALIZATION IN ROBOTICS AND MECHATRONICS	1 YEAR	POST-GRADUATE		X	
PERU	TECHNOLOGICAL UNIVERSITY OF PERU (UTP)	MECHATRONIC ENGINEERING	5 YEARS	UNDER-GRADUATE		X	
	PONTIFICAL CATHOLIC UNIVERSITY OF PERU (PUCP)	MECHATRONIC ENGINEERING	5 YEARS	UNDER-GRADUATE		X	
	PRIVATE UNIVERSITY OF THE NORTH (UPN)	MECHATRONIC ENGINEERING	5 YEARS	UNDER-GRADUATE		X	
	UNIVERSITY OF ENGINEERING AND TECHNOLOGY (UTEC)	MECHATRONIC ENGINEERING	5 YEARS	UNDER-GRADUATE		X	
MEXICO	TECHNOLOGICAL UNIVERSITY OF MEXICO (UNITEC)	MECHATRONIC ENGINEERING	5 YEARS	UNDER-GRADUATE		X	
	NATIONAL AUTONOMOUS UNIVERSITY OF MEXICO (UNAM)	MECHATRONIC ENGINEERING	5 YEARS	UNDER-GRADUATE	X		
	VALLE DE MÉXICO UNIVERSITY (UVM)	MECHATRONIC ENGINEERING	9 SEMES-TERS	UNDER-GRADUATE		X	
	POLYTECHNIC UNIVERSITY OF GUANAJUATO (UPG)	ROBOTICS ENGINEERING	N/D	UNDERGRADU-ATE AND POSTGRADUATE	X		
ARGENTINA	NATIONAL UNIVERSITY OF CUYO (UNCUYO)	MECHATRONIC ENGINEERING	5 YEARS	UNDER-GRADUATE	X		
	UNIVERSITY OF BUENOS AIRES (UBA)	INDUSTRIAL AUTOMATION	1.5 YEAR	POSTUNDER-GRADUATE	X		
	NATIONAL UNIVERSITY OF LA MATANZA (UNLAM)	INDUSTRIAL AUTOMATION	5 YEARS	UNDER-GRADUATE	X		UNIVERSITY-LEVEL TECHNICAL QUALIFICATION IN ELECTRO-NICS (PARTIAL QUALIFICATION) TELECOMMUNICATIONS AND TELEMATICS, ROBOTICS, MULTIMEDIA (MAJORS)
	NATIONAL TECHNOLOGICAL UNIVERSITY (UTN)	AUTOMATION AND ROBOTICS	3 YEARS	UNDER-GRADUATE	X		TECHNICAL QUALIFICATION HIGHER-LEVEL TECHNICAL QUALIFICATION IN AUTOMA-TION AND ROBOTICS
	PONTIFICAL CATHOLIC UNIVERSITY OF ARGENTINA (UCA)	ELECTRONIC ENGINEERING	5 YEARS	POST-GRADUATE		X	DEGREE IN ARTIFICIAL SUPERINTELLIGENCE
	BUENOS AIRES INSTITUTE OF TECHNOLOGY (ITBA)	AUTOMATION AND ROBOTICS	5 YEARS	UNDER-GRADUATE		X	MAJOR IN MECHATRONICS
JAPAN	KYOTO UNIVERSITY	MECHATRONICS	N/D	UNDER-GRADUATE	X		
	NAGOYA UNIVERSITY (CENTER FOR MICRO-NANO MECHATRONICS)	NANOMATERIALS SCIENCE	N/D	POST-GRADUATE	X		ARTIFICIAL INTELLIGENCE APPLICATIONS
		NANODESIGN AND MANUFACTURING	N/D	POST-GRADUATE	X		ARTIFICIAL INTELLIGENCE APPLICATIONS
		NANOMEASUREMENT ENGINEERING	N/D	POST-GRADUATE	X		ARTIFICIAL INTELLIGENCE APPLICATIONS
		NANOCONTROL ENGINEERING	N/D	POST-GRADUATE	X		ARTIFICIAL INTELLIGENCE APPLICATIONS
	TOHOKU UNIVERSITY	ELECTRONIC ENGINEERING	N/D	UNDER-GRADUATE	X		

COUNTRY	UNIVERSITY	DEGREE PROGRAM	DURATION	LEVEL	PUBLIC	PRIVATE	OBSERVATIONS
GERMANY	UNIVERSITY OF STUTTGART	MECHATRONICS	6 SEMESTERS	UNDERGRADUATE	X		UNDERGRADUATE: 6 SEMESTERS, 4 BASIC AND 2 SPECIA-LIZATION; M.SC.: 4 SEMESTERS, 12-WEEK INTERNSHIP, AND 5 SPECIALIZATION OPTIONS
	UNIVERSITY OF ERLANGEN-NUREMBERG (FAU)	MECHATRONICS	N/D	UNDER-GRADUATE	X		
	UNIVERSITY OF HANNOVER	MECHATRONICS	2 YEARS	M.SC.	X		
	RAVENSBURG-WEINGARTEN UNIVERSITY OF APPLIED SCIENCES	MECHATRONICS	3 SEMESTERS	M.SC.	X		
FINLAND	TAMPERE UNIVERSITY OF TECHNOLOGY	M.SC. IN ROBOTICS AND ARTIFICIAL INTELLIGENCE	N/D	POST-GRADUATE	X		CREDIT SYSTEM
	SAIMAA UNIVERSITY OF APPLIED SCIENCES	DEGREE PROGRAM IN ME-CHANICAL ENGINEERING AND PRODUCTION TECHNOLOGY	3,5 YEARS	UNDER-GRADUATE	X		
UNITED KINGDOM	MIDDLESEX UNIVERSITY	M.SC. IN MECHATRONIC SYS-TEMS ENGINEERING	2 YEARS	POST-GRADUATE	X		
UNITED STATES	UNIVERSITY OF MARYLAND	M.ENG. ROBOTICS	N/D	POST-GRADUATE	X		
	WESTERN ILLINOIS UNIVERSITY (WIU)	MECHANICAL ENGINEERING	N/D	UNDER-GRADUATE	X		MAJOR IN ROBOTIC ENGINEERING
	TENNESSEE STATE UNIVERSITY (TSU)	PH.D. IN COMPUTER AND INFORMA-TION SYSTEMS ENGINEERING WITH A CONCENTRATION IN ROBOTICS	N/D	POST-GRADUATE		X	
		PH.D. IN ROBOTICS AND COMPUTER INTEGRATED MANUFACTURING	N/D	POST-GRADUATE		X	
	SANTA CLARA UNIVERSITY	M.SC. IN COMPUTER SCIENCE AND ENGINEERING	5 YEARS	POST-GRADUATE		X	
	OREGON STATE UNIVERSITY, COLLEGE OF ENGINEERING, COLLABO-RATIVE ROBOTICS AND INTELLIGENCE SYSTEMS (CORIS) INSTITUTE	MASTER'S IN ROBOTICS	N/D	POST-GRADUATE		X	
	UNIVERSITY OF GEORGIA	ELECTRICAL AND ELECTRO-NICS ENGINEERING	4 YEARS	UNDER-GRADUATE		X	MAJOR IN ROBOTICS
	MASSACHUSETTS INSTITUTE OF TECHNOLOGY (MIT)	ELECTRICAL ENGINEERING AND COMPUTER SCIENCE	N/D	UNDER-GRADUATE AND POST-GRADUATE		X	ALSO HOME TO THE COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABO-RATORY (CSAIL)
CANADA	UNIVERSITY OF WATERLOO	MECHANICAL ENGINEERING	4 YEARS	UNDER-GRADUATE		X	THE MOST INNOVA-TIVE PROGRAM IN CANADA, KNOWN AS "THE NEW HARVARD"
	POLYTECHNIQUE MONTRÉAL (UDEM	MECHANICAL ENGINEERING	CREDIT SYSTEM	POST-GRADUATE	X		SPECIALIZATION IN AUTOMATION AND SYSTEMS, ROBOTICS AND APPLIED INFOR-MATION
ISRAEL	BEN-GURION UNIVERSITY OF THE NEGEV (BGU)	M.SC. IN ELECTRICAL AND COMPUTER ENGINEERING	2 YEARS	POST-GRADUATE	X		WORLD-RENOWNED RESEARCH CENTER
	TEL AVIV UNIVERSITY (TAU)	AI MACHINE LEARNING	N/D	UNDER-GRADUATE	X		HOME TO THE YANDEX INITIATIVE IN MACHINE LEARNING, A COLLABO-RATION BETWEEN THE BLAVATNIK SCHOOL OF COMPUTER SCIENCE (TAU) AND YANDEX (RUSSIA)

Note: These institutions are ranked among the best in the world by the Academic Ranking of World Universities (ARWU or Shanghai Ranking).
Source: Compiled by the author.

20,000
SCHOOLS IN BRAZIL ARE USING GEEKIE

ing conventional methods. It should be mentioned that the schools that use these systems are private establishments in urban areas attended by upper-middle-class students, whose families can afford to pay reasonably high fees.

Approximately 500,000 third- to ninth-grade students at 590 of India's 3,495 private schools were assessed between 2010 and 2014.[2]

Brazil's Geekie platform, when used for four hours per day for eight weeks, helps students attain grades that are 30% higher, on average, than those of young people who have not used it. It is being used at over 20,000 schools by 5 million students in Brazil.[3]

Incentives for Curiosity

Adaptive learning and instructional models for teachers to design contents are just some of the new education tools available. This prompts the question of how AI will impact teaching as a whole.

One aspect of is the use of adaptive learning models. The use of algorithms in the classroom could help detect teaching deficits and give teachers real-time feedback. This could lead to personalized digital tutoring that would improve the learning experience.

Another area of application of these models are adaptive assessments. Questions are selected based on students' responses to earlier test questions, a score is estimated, and then their performances are analyzed. This data collection process generates a continuous feedback cycle that can bring about major improvements in the learning process. This data would provide the teacher with information on the student's academic performance in different areas and would thus generate more holistic assessments of this.

How will the role of teachers change with the advent of these new technologies? Are their jobs at risk? Some of the more apocalyptic views on this issue conceal a fallacy. The teacher is and will remain a key figure in the learning process and their critical perspective will continue to be essential, promoting discussion among students, the exchange of personal experiences, social interaction, collaborative work, and research.

Justin Reich (2014) analyzed computer assessment to consider where this trend is headed. He argues that computers that have not been trained by humans are generally very good at assessing quantitative factors. In other words, they stand out when carrying out assessments that humans no longer need to perform and function as a support system, rather than as competition.

New technologies will be able to generate more individualized texts, depending on each student's needs. This is why McGraw-Hill, Pearson, Santillana, and other publishing houses have been developing new adaptive technology tools that understand students and anticipate what they do not know. By using this information, the system provides personalized, algorithmic contents that are tailored to individual needs. In other words, what we are witnessing is a much more developed form of involvement than mere quantitative analysis.

30%
THE INCREASE IN STUDENTS' GRADES WHEN THEY USED DIGITAL LEARNING PLATFORMS

Martín Molina, a professor at the Artificial Intelligence Department at the Technical University of Madrid, AI and robotics are still in the early stages. No robots yet have visual or linguistic comprehension abilities: "They cannot understand an entire book or translate a poem, as they are unable to truly comprehend, let alone summarize, a text," he argues (Schapira, 2017).

The predisposition and ability to engage in social interactions, leadership, and teamwork, all of which are particularly human traits, are also hard to emulate through robotization and AI.

How can we program curiosity, which is what prompts us to research things, study them, and fill in the holes in our information? What factors motivate individuals to cross the barriers of their own knowledge and experiences? To paraphrase Santiago Bilinkis (2017), might this be our secret weapon against the robots?

These changes are challenging education policy-makers in new ways. In the transition toward widespread use of these innovations, efforts need to be directed toward generating regulations for standardizing the use of these applications in the classroom, improving teacher training, protecting the privacy of student data, and allowing interinstitutional articulations that help these technologies coexist with traditional systems.

AI will certainly disrupt education, but there is a long road ahead. The challenges revolve around finding ways to capitalize on the benefits of a more personalized form of teaching, use new technologies to expand and open up the classroom, articulate face-to-face and online learning, and increase interactions between teachers and students.

NOTES

[1]One of the top 50 US universities.
[2]IE has implemented Mindspark learning support centers in New Delhi since 2013 and has also focused its strategy on vulnerable communities.
[3]See ProFuturo (2016).

REFERENCES

Barrat, J. 2013. *Our Final Invention: Artificial Intelligence and the End of the Human Era*. New York: Thomas Dunne Books.
Bilinkis, S. 2017. "El arma secreta de los humanos contra los robots." *La Nación.* July 30.
Cataldi, Z., and Lage, F. J. 2010. "Modelado del estudiante en sistemas tutores inteligentes." *TE&ET, Revista Iberoamericana de Tecnología en Educación y Educación en Tecnología* 5: 29–38.
Escamilla, J., Calleja, B., Villalba, E., et al. 2014. "Aprendizaje y evaluación adaptativos." Edu Trends report, July. Observatorio de Innovación Educativa del Tecnológico de Monterrey.
Gross, B. 1992. "La inteligencia artificial y su aplicación en la enseñanza." *Comunicación Lenguaje y Educación* 13: 73–80.
Kaplan, J. 2016. *Artificial Intelligence: What Everyone Needs to Know*. New York: Oxford University Press.
Morillo Lozano, M.C. 2016. *Aprendizaje adaptativo*. Master's thesis on technology and information technology. Valladolid: Universidad de Valladolid.
ProFuturo. 2016. "Geekie: aprendizaje personalizado para todos." *ProFuturo*. December 20.
Reich, J. 2014. "Computers Can Assess What Computers Do Best." Education Week April 29.
Salgueiro, F.A. 2005. *Sistemas Inteligentes para el modelado del tutor*. Undergraduate thesis on IT Engineering. Buenos Aires: University of Buenos Aires.
Schapira, D. 2017. "La inteligencia artificial: una revolución de la educación." *El Cronista*. September 25.
Walsh, T. 2018

INTERVIEW

Rafael Reif
President of
Massachusetts Institute
of Technology (MIT)

"Every technological project has an ethical dimension to it"

VENEZUELAN-BORN RAFAEL REIF LEADS THE MASSACHUSETTS INSTITUTE OF TECHNOLOGY (MIT), ONE OF THE WORLD'S LEADING EDUCATIONAL AND RESEARCH CENTERS. IN THIS INTERVIEW WITH *INTEGRATION & TRADE*, REIF HIGHLIGHTS THE ROLE OF ARTIFICIAL INTELLIGENCE AS A COMPLEMENT TO HUMAN INTELLIGENCE, POINTS OUT THE RISKS OF JOB AUTOMATION THAT IT POSES AND HOW TO TACKLE THESE, AND EMPHASIZES THE NEED TO INCLUDE ETHICAL SPECIALISTS DURING THE DESIGN STAGES FOR TECHNOLOGICAL PROJECTS.

What are the main challenges facing research into artificial intelligence (AI)?

Although it sometimes feels like AI is everywhere we look, many people might not know a few underlying truths about it. One is that the foundations of AI in use right now are, in fact, relatively old. There are teams of brilliant researchers at our biggest companies that, with few exceptions, focus mostly on trying to squeeze new applications out of existing approaches. If we want to achieve the kind of breakthroughs that will revolutionize the field, it's going to take new science. Just imagine if the next breakthrough in AI comes from the root of intelligence itself: the human brain. That's how big we need to think to generate new knowledge in this area, as well as enable its practical impact.

How do you tackle this challenge?

To approach this challenge, we have launched the MIT Quest for Intelligence, an MIT-wide initiative on human and machine intelligence. Its aim is partly to increase our understanding of human intelligence, and partly to drive the development of technological tools that we believe can positively influence virtually every aspect of society. We are looking for the next great leap that will unlock innovations that today's entrepreneurs worldwide haven't begun to imagine. I find that very exciting.

> IT IS CRITICAL FOR TECHNOLOGISTS TO INVITE THEIR COLLEAGUES TO COLLABORATE

Why might this technology be important for Latin America and the Caribbean?

These tools have the potential to help us address a number of complex challenges across society, in Latin America and beyond. The breadth of application is staggering, with relevance in the areas you mention, but also in fields like human health, manufacturing, finance, and many more. In terms of poverty, I think about the potential impact for an entity like MIT's Abdul Latif Jameel Poverty Action Lab, or J-PAL. Its mission is simple: to reduce global poverty. To determine which poverty programs

LEO RAFAEL REIF GROISMAN WAS BORN IN MARACAIBO, VENEZUELA, IN 1950. HE RECEIVED HIS UNDERGRADUATE DEGREE IN ELECTRICAL ENGINEERING THERE. AFTER EARNING A DOCTORATE AT STANFORD, HE BEGAN HIS TEACHING CAREER AT MIT IN THE EARLY 1980S. HIS RESEARCH FOCUS WAS THREE-DIMENSIONAL INTEGRATED CIRCUIT TECHNOLOGIES AND ENVIRONMENTALLY FRIENDLY MICROELECTRONICS FABRICATION. HE ALSO WORKED TO CREATE MITX, AN ONLINE EDUCATION PROJECT DESIGNED TO PROVIDE MIT STUDENTS WITH NEW TOOLS AND TO MAKE MIT COURSES FREELY AVAILABLE TO STUDENTS AROUND THE WORLD. AFTER 25 YEARS ON THE FACULTY AT MIT, HE WAS APPOINTED PRESIDENT IN 2005.

are actually making a difference in the US and around the world, its researchers test the programs through double-blind randomized evaluations. If those evaluations could be enhanced with new AI tools, I think the lab could make important gains. We are just scratching the surface of what is possible. I am hopeful that leveraging human and machine intelligence will generate meaningful results for all of society.

THE ETHICAL QUESTIONS SURROUNDING AI MUST BE A CENTRAL CONCERN.

Is a bias-free form of artificial intelligence possible?

This is an interesting question, and one I'm not sure I'm qualified to answer. But there are sociologists and humanists and ethicists at MIT and beyond who are working on this. For instance, there is a project at the MIT Media Lab called the Algorithmic Justice League, which is committed to fairness, accountability, and transparency in coded systems. It is critical that technologists invite their colleagues to collaborate at the start of the design phase, not once a product is nearing completion. These are important issues that will require the wisdom of many early in the process.

What are the next steps in this direction?

One solution is to always encourage participation from across a broad range of backgrounds and perspectives. The more we can bring together researchers who reflect the diversity of the human experience, the better our chances of minimizing bias. That's true with AI. And frankly, I believe it's true in anything we do.

LATIN AMERICA HAS ENORMOUS INTELLECTUAL CAPITAL WITH WHICH TO SOLVE LOCAL PROBLEMS

How will artificial intelligence impact the world of work?

AI does have the potential to generate exciting new opportunities in productivity and economic growth, especially in nations that are able and willing to make an investment in it. My concern is that it also has the potential to become a source of inequality, as well, both between nations and within them. That makes me uneasy. I firmly believe that getting this right is among the most important and inspiring challenges of our time. The ethical questions surrounding AI must be a central concern, not an afterthought. How can we design these technologies from the start to serve the best interests of our whole society?

I AM CONCERNED THAT TECHNOLOGICAL PROGRESS WILL BECOME A SOURCE OF INEQUALITY

I believe that the answer to that question is not simple. Do you have any thoughts on it?

We are thinking hard about that question at MIT. We're looking, in particular, at the impact technology is having on jobs. We have all seen the influence automation has on our work, our lives, and our society. At MIT, where researchers shape technologies that will shape the work of tomorrow, I believe we have a special responsibility to consider the consequences of those technologies. Over the winter, we launched a two-year study that we are calling the MIT Task Force on the Work of the Future. This project aims to help us understand how technologies are transforming the nature of work and identify strategies to shape technological innovation so that it complements, rather than replaces, human workers. The world is changing, and so are jobs. There will, of course, be work in the future, but what those jobs look like will be up to us.

Can you give us an example where research into AI has contributed to improving human well-being?

The stakes are so high that we must find ways to collaborate with colleagues from around the world in finding solutions to really difficult problems. The power of AI will come from applying these new tools to really hard problems that we haven't been able to solve until now. Here's an example of what I mean. Regina Barzilay, an MIT professor of electrical engineering and computer science, studies machine learning and uses it to do all kinds of remarkable things, like decode ancient languages. But in 2014, when she was diagnosed with breast cancer, she saw that although clinical oncology produces massive amounts of data about patients and the outcomes of their treatments, the data is severely underutilized. Desperate for facts, she set out to find them using machine learning. Professor Barzilay's dream is to leverage the extraordinary promise of machine intelligence to revolutionize cancer care—to use a treasure trove of information to identify patterns and personalize treatment. In fact, that dream is already becoming reality: she and her research team recently deployed machine learning tools to several area hospitals.

CASE STUDY

Closing the Digital Divide

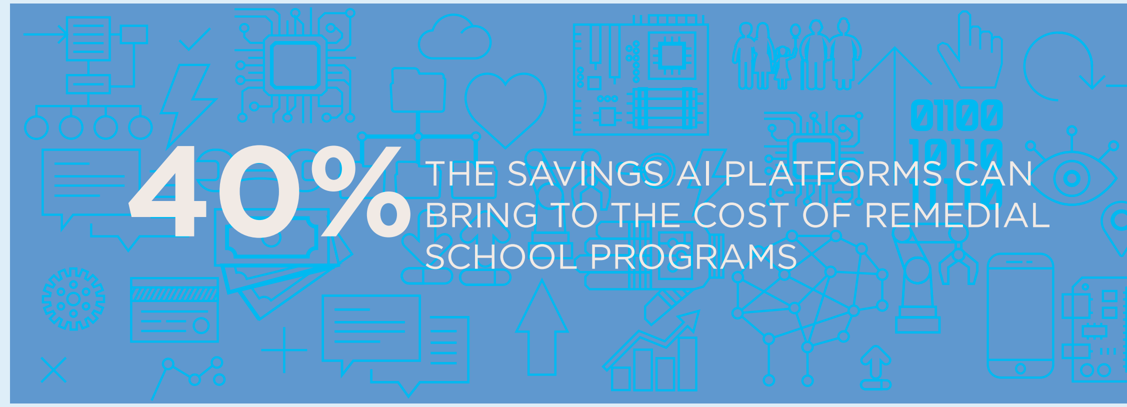

Brainy is an artificial intelligence-based tool that is being used in the Chilean education system. It was created by the firm Cognitiva in partnership with Red Crecemos and provides support for teachers and offers complementary help with planning activities that are tailored to each student's needs. Cognitiva's founder and CEO Rolando Castro emphasizes the importance of AI in closing the digital divide in formal education.

Why are you seeking to include AI in education?

Our mission is to improve Latin Americans' quality of life and contribute to democratizing knowledge. Projects in which we can use AI to bring children in vulnerable situations closer to schools and facilitate their access to education is in line with that mission. Technology today doesn't just enable access to content in an exponential fashion, it also improves the way students engage with this content, deepening their understanding of their school subjects through online games, illustrative videos, and interactive images that make learning memorable for children from an early age.

In Chile, you have developed a digital assistant called Brainy in partnership with Red Crecemos. What does this tool do?

Brainy is the first concrete AI-based solution that has been applied to education in Chile. It provides support for teachers and also helps them plan activities that are tailored to each student's needs so that they can learn better, faster, and develop different life skills. This is an unprecedented initiative for both Chile and Latin America and it is targeting schools attended by children from vulnerable social contexts. This educational platform has led to interesting scientific findings and seeks to help students get better qualifications and increase their motivation to learn. It is designed for children between the ages of 9 and 14 and understands questions they ask it in everyday Chilean Spanish, which it then provides age-appropriate answers to. Brainy will bring down the cost of Chile's remedial school programs by 40% and increase access to specialized content and personalized learning by 25%.

How many students are using it?

Brainy aims to improve academic excellence in schools in vulnerable social contexts. It focuses not just on academic performance terms, but also by supporting and developing students' soft skills. It seeks to bridge the educational divides that currently exist in Chile and improve vulnerable students' access to knowledge. It began to be implemented in March 2018 in fourth-grade science classes at three Red Crecemos schools. Over the next four years, other grades and subjects will be added to the program, which will be extended into three more schools, reaching a total of approximately 500 students. The program's real impact won't be confirmed until 2019, when it will have been in operation for a year.

Do you think the program could be scaled up and used more widely?

It is becoming more and more commonplace for teachers to develop innovative approaches to strengthening education and find new ways for their students to build their skills. Brainy was developed to provide support for fourth- and sixth-grade elementary school students in their science and language classes, and we hope that it will be used for other subjects and by other grades at Red Crecemos schools over the next four years. The platform is flexible and complements teachers' work in the classroom as it guides the teaching and learning processes, provides official curriculum content, and establishes a cognitive profile for each student. Interactions are generated within the platform, which connects teachers, tutors, and students using metrics. Brainy it is still at the trial stage—but we hope that by late 2018 it will become a standard fixture in the classrooms of state-subsidized Chilean schools.

Do educational establishments benefit from using the platform?

Yes, provided that schools want to break with traditional educational approaches, improve performances at school by implementing technological systems that involve personalized learning, and bring together all interested parties. Including AI in classrooms means applying innovative methodologies so that students can face the challenges of the future. The inclusion of technology in education translates into a series of learning benefits that spark students' curiosity and motivate them, as they can explore their favorite subjects further and learn through play. These technological solutions involve the entire educational ecosystem. As well as helping students, they provide teachers with educational and administrative support.

Can AI help reduce the social divide and gaps in income and opportunities?

We are convinced that AI can help close the digital divide and democratize information for all. In the field of education, AI can be used to build smarter systems, using the word “smart” to mean the ability to continually adapt to learning environments and different users’ knowledge levels—in this case, those of both students and teachers. Projects may also be generated to favor larger communities and benefit entire regions or geographic areas, providing support for different professionals and supporting the process of decentralizing knowledge.

How far does the platform change classroom dynamics?

Brainy is designed to provide instruction and continuous support for the learning and teaching processes by building, updating, and analyzing factors that reflect each student’s behavior. It is a very advanced pedagogical tool that can deliver individualized learning experiences which help make the teaching and learning processes more flexible and personalized. The fact that it is personalized helps detect which students have more learning difficulties, enabling teachers to focus more on these students and detect possible cases of bullying or other such situations. This effectively contributes to closing the education gap as it helps teachers to identify early on which students need to be given more support and more examples to be able to access content.

What will the education of the future be like?

I believe that the education of the future will be based on the democratization of learning and greater student involvement in the classroom, a radical shift that we are already seeing at some establishments. The evolution of education goes hand-in-hand with the inclusion of technology in different processes, always as a complement to what the teacher does. This system will allow us to take students’ needs into consideration and choose whatever approach to learning is easiest and most relevant to them, as well as incorporating different technologies like virtual reality and AI.

REGIONAL SCOPE

Children can use Brainy in the classroom and at home though a user interface that functions on desktop computers and tablets. It speaks 24 different regional varieties of Spanish, including those of Colombia and Chile. Use of the tool will be extended into the last four years of high school in Chile and it will also be adopted in Peru and other countries in the region.

E-Health

The **Future** of the **Health Sector**

Michael Gillam
HealthLab and Singularity University

GREAT PROGRESS HAS BEEN MADE IN HEALTHCARE IN RECENT YEARS, BUT DESPITE THIS, THERE ARE STILL MANY CHALLENGES TO OVERCOME IF WE ARE TO IMPROVE HUMAN WELL-BEING AND INCREASE LIFE EXPECTANCY. THIS ARTICLE EXAMINES THESE CHALLENGES AND THE OPPORTUNITIES THAT NEW TECHNOLOGIES MAY BRING FOR THE DEVELOPMENT OF KNOWLEDGE-BASED SERVICES (KBSS) IN THE FIELD OF MEDICINE.

The World Health Organization estimates that almost 400 million people lack access to basic and essential health services and two billion lack access to surgical services (Funk et al., 2010). In the US, mistakes in healthcare have become the third leading cause of death behind cancer and cardiovascular disease (Makary and Daniel, 2016). However, new trends in computation have opened up new opportunities for delivering services, including healthcare.

It has been said that the largest taxi company in the world today, Uber, owns no vehicles. The world's largest accommodation provider, Airbnb, owns no real estate. The most valuable retailer, Alibaba, owns no products. Until recently, the world's most popular media company, Facebook, didn't produce any media content. This has led to speculation about what assets the world's largest healthcare company may operate without in the future. Many believe that the answer lies in the growing power of computing, particularly artificial intelligence (AI).

The seeds for that growth were trends that began several decades ago. In 1965, Gordon Moore, the cofounder of Intel, the world's largest PC microchip manufacturer, noticed that every 18 months, his company was able to double the number of components it could fit on a microchip through new advances. Ten years later, Moore revised his estimate to be more precise: this advance took place every 24 months. The decades rolled on and computing power continued to double every two years, virtually like clockwork. One of the fields that would be affected by this was healthcare.

Some 25 years later, in 1990, the Human Genome Project was launched by the National Institutes of Health (NIH), an ambitious US$3 million project to determine the entire sequence of a human genome: 23 pairs chromosomes comprising 3 billion base pairs. There were plenty of critics. One said, "…the Human Genome Project has been sold on hype and glitter…" Dr. Martin Rechsteiner, a biochemist from the University of Utah, said: "…the Human Genome Project is bad science, it's unthought-out science, it's hyped science." Dr. Michael Syvanen, a microbiologist at the University of California at Davis, argued: "everybody I talk to thinks this is an incredibly bad idea" (Angier, 1990).

Eight years into the project the critics appeared to be right. Over half the time was over, but the project was only 4% complete. Scientists needed to finish the remaining 96% of the genome in just seven years. Critics noted that at this rate of progress the project would take another 150 years to finish. But where some saw failure, others saw an opportunity. J. Craig Venter raised hundreds of millions of dollars to create a private company, Celera Genomics, that promised to deliver a full human genome within the remaining time by building one of the largest clusters of computational power in the world.

The project didn't take 150 years to finish. Five years later, not only did the NIH sequence the genome, but Venter's company did too, both publishing a day apart in *Science* magazine. Both finished two years ahead of schedule. Their critics had been wrong in their estimates by over 100 years. What did Venter and others see that the critics had missed? The progress of the Human Genome Project had been exponential rather than linear (see figure 1).

Some authors have attributed our tendency to project trends linearly and the lack of human intuitive ability around exponential trends to our evolutionary heritage. If it takes a deer 30 steps to run across a field, one can intuitively predict where to throw a spear to intercept it (for example, at a point 30 meters away). But if a deer were to take 30 exponential steps (that is, 2 meters, 4 meters, 8 meters, and so on), few would be able to intuit that this would take the deer 26 times around the earth. Good luck to the hunter trying to catch dinner in that sort of world. However, this is the world in which we now live. This is also the world we are increasingly seeing in the healthcare sector, and especially in AI.

NEW CHALLENGES

Whereas the first human genome cost US$2.7 billion in 2001, the latest cost to sequence a human genome is under US$1,000. Remarkably, the cost for sequencing genomes is dropping at a rate five times faster than even Moore's Law (see figure 2). If the current rate of decline continues for five more years, the cost for a full human genome will be less than the cost of ordering a pizza (around US$10). If the trend continues for 10 years, the cost per genome will be less than the cost of flushing a toilet (around US$0.01). This exponential trend portends a world where genomic data may be available ubiquitously across industries inside and outside of healthcare. This rising wave has left many wondering how we will handle, process, and create value with this inundation of genomic data.

FIGURE 1
GROWTH IN DNA CATALOGUED IN GENBANK

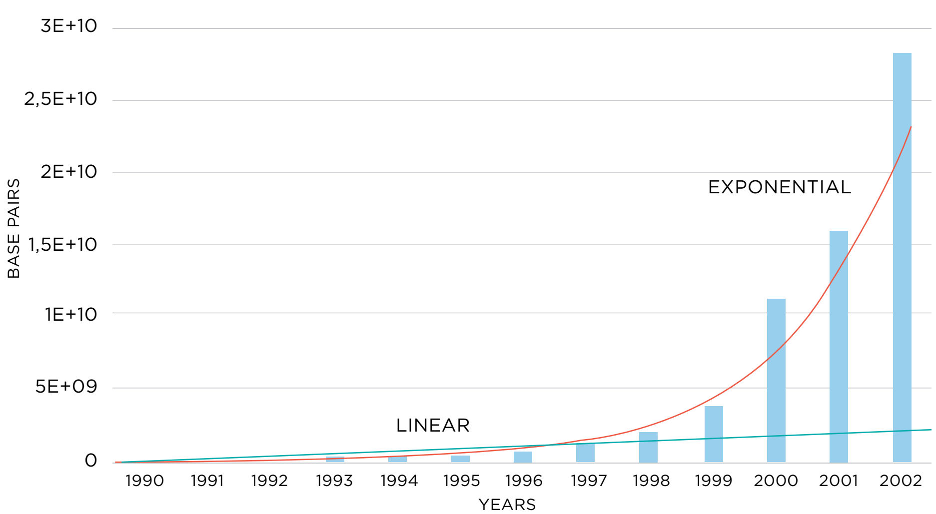

Source: GenBank Statistics, National Center for Biotechnology Information.

FIGURE 2
COST PER HUMAN GENOME

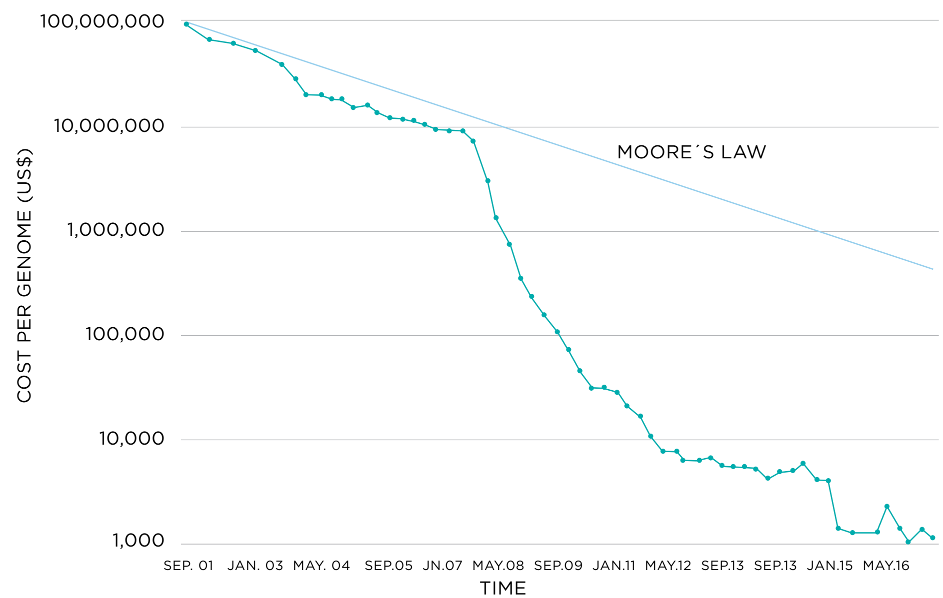

Source: National Human Genome Research Institute.

The exponential rise in genomic data, the complexity of interpreting it, and the unpredictability of modifying it have created new challenges for healthcare. Translating nearly three billion base pair letters of DNA (i.e., A, T, G, and C) into meaningful metabolic pathways (the metabolome) or physical traits (the phenome) remains an unsolved challenge in healthcare. Theoretically, given the genome of an embryo and with enough genomic knowledge, healthcare could predict physical and metabolic traits beyond just hair, eye color, and genetic diseases. Given enough genetic understanding, the metabolic side effects of virtually every medication should be known. The genome of an embryo should enable us to predict the eventual facial appearance of a child right up through adulthood from before the child is even born. Connecting the genome to the metabolome and the phenome remains an unsolved challenge in medicine.

New genomic modification techniques, like CRISPR/Cas9 (CC9), are also facing challenges that seem daunting to the human intellect. CC9 is a search-and-insert (splicing) system that modifies DNA in living cells. CC9 snips open DNA at a location by matching the location to another sequence of DNA it is carrying. After the DNA is opened, CC9 inserts this sequence into the opening. The two sides of the DNA then reattach (reanneal) thus successfully completing the insertion. Though CC9 is extremely accurate in where it opens the DNA, sometimes the reannealing process fails. Often moving the incision point up or down in the genome can result in much more successful CC9 gene modifications but predicting where that splice point might be is guesswork. However, determining successful splicing points is the key to unlocking many successful gene modifications.

These two healthcare challenges—translating the meaning of the genome and modifying the genome—are being tackled by new approaches in AI, particularly deep learning.

DEEP LEARNING

In 2012, years after early efforts in AI that attempted to imitate the neural structure of the brain had stalled, Geoffrey Hinton's lab showed a dramatic improvement in image recognition accuracy at the ImageNet image recognition competition (see figure 3). Hinton called their approach "deep learning." It was just a few years later, through a series of improvements to deep learning, that computers began to show lower error rates than humans (around 5%) at recognizing images in the competition. It was in in 2016 that computers surpassed human ability to recognize images from the ImageNet dataset. Subsequent groups have continued to extend that lead, achieving deep learning image recognition error rates less than half those of humans.

After Hinton's demonstration, progress in deep learning began spanning multiple domains of AI. In 2016, Microsoft announced they had built a deep learning system that can translate audio text from cell phones in a way that exceeds human ability (Xiong et al., 2016). Importantly, in contrast to static images, their approach demonstrated that advances could be made with data streams with a temporal element.

Today, AI efforts, including both deep learning and other areas, have led to progress across vastly heterogeneous tasks. The Electronic Freedom Foundation (EFF) tracks the metrics for progress in machine learning compared to human rates across scores of tasks.[1]This progress has begun to spill over into healthcare.

FIGURE 3
IMAGENET IMAGE RECOGNITION CONTEST ERROR RATES

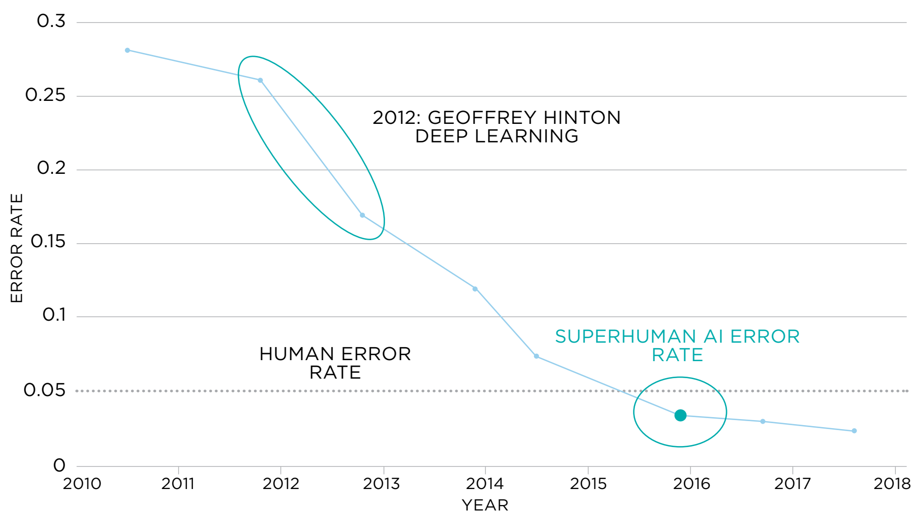

Source: Approximated from https://www.eff.org/ai/metrics

FIGURE 4
NUMBERS OF DEEP LEARNING ARTICLES IN THE MEDICAL LITERATURE PER YEAR

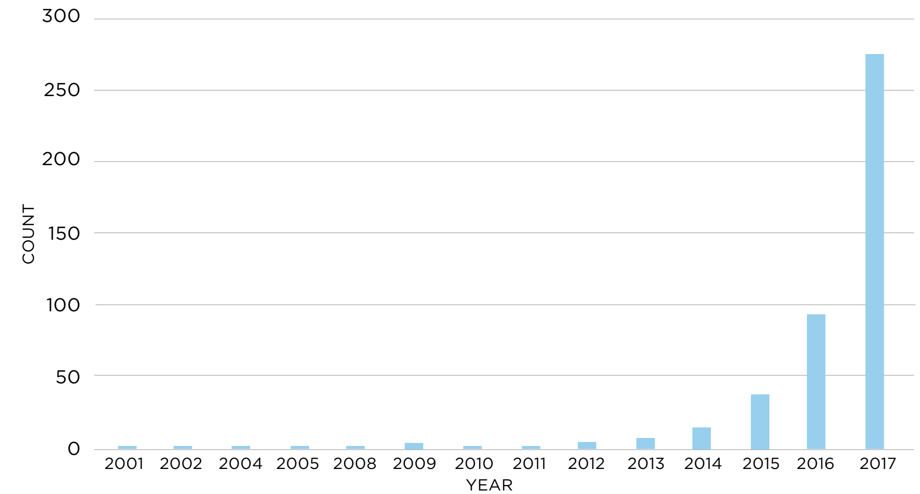

Source: Author's calculations based on PubMed.gov (2017)

DEEP LEARNING IN HEALTHCARE

The revolution in the ability of AI to recognize images has led to a stream of advances in healthcare services. In November 2016, Google partnered with Britain's National Health Service (NHS) to gain access to ophthalmologic images. Just six months later, Google announced they had built a deep learning system that could recognize damage to the retina from diabetes (diabetic retinopathy) from images with a skill that matched board-certified ophthalmologists (Peng and Gulshan, 2016).

In the summer of 2017, Stanford published on a system that could classify pictures of skin cancer as benign or malignant as effectively as board-certified dermatologists (Kubota, 2017).

Jeremy Howard, the founder of the data scientist contest platform Kaggle.com and CEO of Enlitic, has announced that his company could diagnose wrist fractures with an area under the curve (AUC) of 0.97 ("more than three times better than the 0.85 AUC achieved by leading radiologists") and could read "chest CT images 50% more accurately than an expert panel of thoracic radiologists."[2] His company has deployed a system for live clinical use which moves any CT scans it sees with cerebral hemorrhages to the front of the triage line of studies to be read by radiologists. Howard noted this month that he is no longer surprised by any studies showing computers exceeding human ability in classifying or recognizing features from images. He added that it is time to just assume that all images can be read by computers better than humans and that the real challenge in medicine is the adoption of these technologies.[3]

Deep learning progress in healthcare has extended far beyond image data. In the summer of 2017, Stanford published a study describing a system that could diagnose 14 different cardiac rhythms from rhythm strips as accurately as cardiologists could. The system was trained on data using the Zio iRhythm, a portable two-lead electrocardiogram monitor. The manufacturer of the Kardia Band, a two-lead wearable ECG for the Apple Watch which has recently received FDA approval, announced they are adding neural networks to interpret the rhythms collected from the device.

The problem mentioned above of determining the best splicing sites for genes is now also being tackled using deep learning. Companies like Deep Genomics are now using this approach to take on the challenge of transforming gene sequences into the metabolome (and eventually the phenome). Deep learning is also being used to increase the accuracy of the CRISPR-Cas9 targeting systems mentioned above, thus making gene modification more successful (Doyle, 2016).

One remarkable aspect of genomic progress is that the systems make their predictions with a high degree of complexity. Healthcare is reaching a moment when AI systems will make discoveries and predictions but may never be able to explain their understanding to us because their complexity exceeds the capacity of the human mind. The black box nature of many AI systems remains a challenge if AI is to achieve the trust levels that could facilitate more rapid adoption in the healthcare sector.

Publications cataloged in PubMed.gov as having "deep learning" in their title have increased exponentially in the healthcare sector (see figure 4). For the moment, the growth curve shows no signs of leveling off.

Other AI technologies apart from deep learning have also gained traction. Human behavior has long been thought to be so complex as to defy accurate prediction. Some AI systems appear to be making striking progress in predicting human behavior. In the fall of 2017, Vanderbilt described a system able to predict suicide attempts within two weeks with 92% accuracy and with 80%–

FIGURE 5
EXPONENTIAL GROWTH OF COMPUTING AVAILABLE PER US$1,000

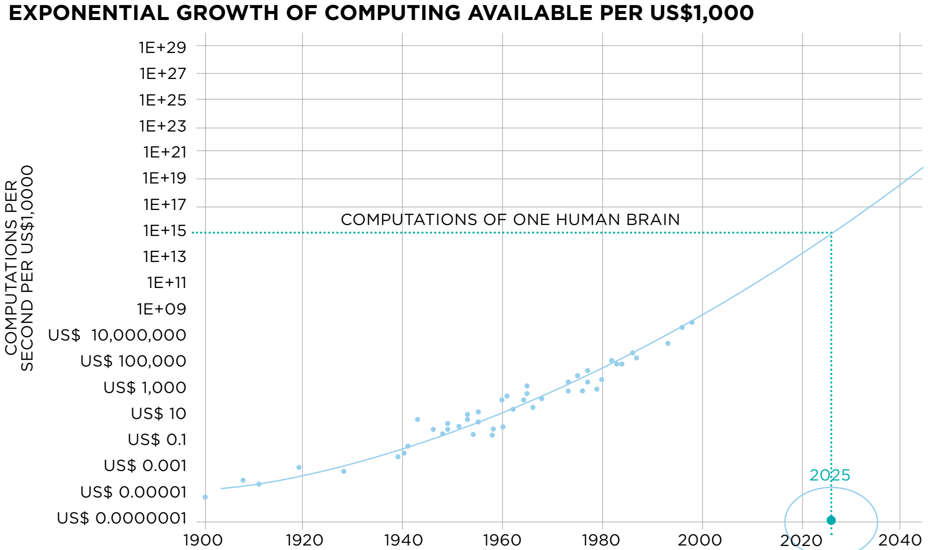

Source: http://www.singularity.com/charts/, page 17.

FIGURE 6
NEURAL SIMULATION SPEED

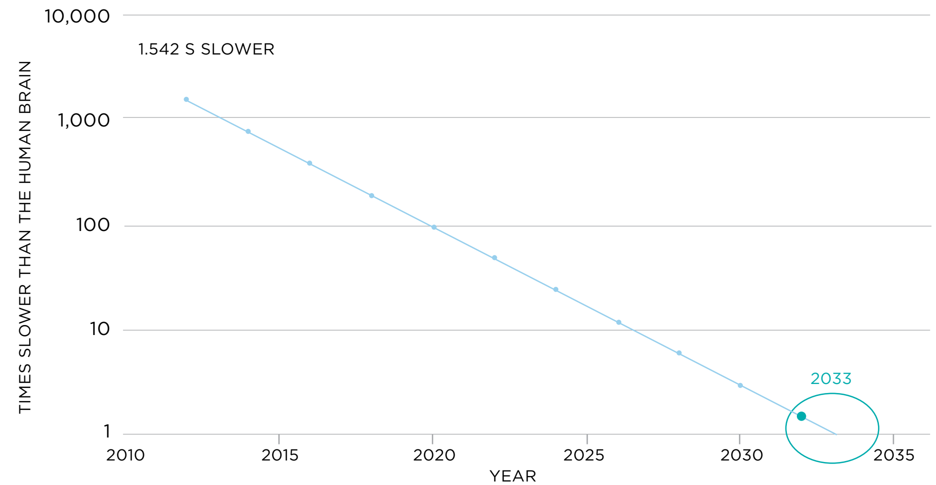

Source: IBM Research (2012).

90% accuracy within two years. This system leveraged a machine intelligence approach called random forests.

ACCESS TO INFORMATION

Apart from the black box challenge of understanding how AI is making its decisions, there are a variety of other challenges facing the development of AI systems for healthcare.

Training AIs in healthcare is a challenge due to a lack of simulated environments. In October 2017, Google's DeepMind team demonstrated a new version of their AlphaGo system (Alpha Go Zero) that learned superhuman skills at the Korean game of Go by playing itself over and over (Vincent, 2017a). Whereas the first version of AlphaGo, which beat 18-time world champion Lee Se-dol, relied upon 100,000 human games as a starting dataset, this new version took just three days of playing itself to beat that first version 100 games to 0. Descendants of this system went on to beat the world's best chess-playing program, Stockfish, after playing chess against itself for four hours, then beat the world's best Shogi-playing program, Elmo, after just two hours of playing itself. These games had the shared feature that practicing could be done in a completely virtual world. In 2016, both Google and Elon Musk's OpenAI released AI sandbox playgrounds for training AI systems (Burgess, 2016).

The healthcare sector lacks high-fidelity simulations of metabolic and physiological processes that could be used by AI for the multiple iterations often needed. Cleansing data of bias is another challenge. AI has been shown to learn biases inherent in the data sets used for training (Hsu, 2017). Racial and ethnicity biases have also been found in diagnostic and treatment data in healthcare (Vedantam, 2007). Challenges remain around how to train AI systems without biases when the training data itself has unknown or unrecognized racial or ethnic biases. The healthcare sector has yet to answer the question of how we build AI to be better than ourselves.

Healthcare datasets are particularly known for their temporality and sparsity. Today, the greatest progress seen in images and audio datasets are for datasets that are complete: pixel data is present for every pixel in the image or audio data is available for every time step. In the US, there are over 60,000 CLIA-certified lab tests and over 1,000 more such tests are added every year. However, not every patient has every lab test performed on them and few lab tests are done on a fixed timeline. AI systems currently face challenges when dealing with the combination of temporality and sparsity. Research groups such as at the University of California are working to make progress in these areas in the healthcare sector.[4]

Access to patient data can present another barrier to AI progress. Sharing private healthcare data across health systems can be challenged by business interests between competitive healthcare systems or patient privacy interests.

However, there may be emerging solutions for data sharing that preserve privacy. Typical cryptographic hashing techniques leave data in a form in which even simple math becomes impossible, let alone sophisticated AI learning. For example, a single lab value like “15” can become a 40-digit alphanumeric string when cryptographically hashed using SHA256. New forms of cryptographic hashing, particularly homomorphic encryption, hold some promise as ways to encrypt data while still allowing computation to be performed. The company Numerai built an entire hedge fund in a process where their data was shared to data scientists in an encrypted format where computation could still be performed by data scientists. The data remains essentially unknown to the data scientists and the company, which is paid in bitcoin (Metz, 2016). Homomorphic encryption holds promise for enabling shared computation on healthcare data in ways that preserve patient privacy.

FIGURE 7
BRAIN SCANNING RESOLUTION

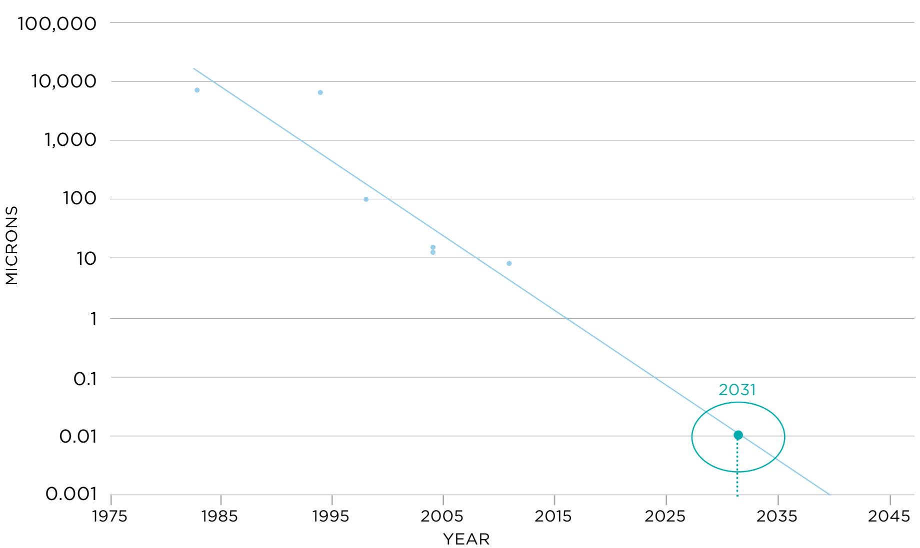

Source: Kurzweil (2014).

THE FATE OF AI IN HEALTHCARE

Despite the barriers, three trends are driving optimism for the increasing role that AI can play in healthcare: the exponential rise in computational power, simulation capacity, and parallelized computation models. These three trends seem to suggest that there may soon come a time when computation will reach parity with the human brain and thus more of the varied capabilities exhibited by human intellect will become available as computational services.

The human brain has roughly 86 billion neurons and an average 10,000 connections (synapses) between each neuron. Computation is performed at the synapses. The map of 1,000 trillion (10^15) connections is called the connectome. Assuming a single calculation per synapse, if current trends in the number of digital-calculations-per-second available at a cost of US$1,000 continue, then computers that could reach computational parity with the human mind will be available by 2025. To put this trend into perspective, if we take the iPhone X, the tenth-anniversary edition of the iPhone, which was released in 2017, as a starting point, the potential iPhone XX, would be released in 2027, could perform calculations at a higher rate than the brain of the person buying it.

Reaching computational parity with the human brain does not indicate that digital designs will exhibit the same parallel processing capability as mammalian brains. Estimates from neurological simulations are needed for this.

Simulation Capacity

Designing computer chips that simulate neural structure (neuromorphic computing) suggests a similar timeline for human brain parity as pure computation. In 2012, as part of the SyNAPSE DARPA project, IBM built a system to simulate the connections of a human brain. The system consisted of roughly six times as many neurons as a typical human brain (530 billion) and about one-tenth the number of connections (100 trillion). The resulting simulation was 1,542 times slower than an actual human brain (IBM Research, 2012). Projections for when this neuromorphic simulation might run at parity with human neurological capacity, without accounting for any improvements in software or hardware design, put the year at 2033. Given potential design improvements, this could arrive earlier.

Capturing the Entire Connectome

Though we may have the computational power and parallelism needed to simulate a human mind, the correct wiring diagram would still be needed. For this, we need the connectome. The challenge with capturing a connectome is that when water freezes, it expands. The result is that during cryopreservation of tissues, the synapse is typically disrupted due to water expansion, meaning that the connections and functional knowledge are lost. This was the case until February 2016, when the Brain Preservation Foundation’s Large Mammal Brain Preservation Prize was awarded. The winning project opened up the possibility of freezing virtually any species’ brain while preserving the connectome (The Brain Preservation Foundation, 2017). This means that it will now be possible to scan human brains for their connections, which can then be represented digitally. Accounting for recent studies suggesting that the width of the synaptic connection’s endplate is important, current scanning technology would need to capture resolutions at 0.01 microns in width. Current trends in destructive scanning resolution capability suggest that the human connectome could be captured at that resolution by 2031 (see figure 7). Some have suggested that this trend suggests a time where even if we don't fully understand how the brain works, we may be able to digitally instantiate full connectomes that operate similarly to human cognition (Kurzweil, 2014). While AI moves toward this point, it is expected that it will also unlock additional capabilities of the human mind as scanning resolution captures deeper structures of the connectome.

Taking these three trends together with their end dates—computation: 2025, simulation: 2033, and scanning resolution: 2031—has led to some suggesting that the computation available for AI in healthcare and other fields will reach human mind parity around 2029 (Kurzweil, 2014).

Though this approach may underestimate the actual number of computations of the human brain, the trend is suggestive of the capacity of AI ahead unlocked by exponential trends.

Brain Control Interface Bandwidth

Some believe that AI progress will outpace the human mind, while others believe advances will occur in tandem. In 2011, Cathy, a 59-year-old mother of two who had been paralyzed from the neck down by a brainstem stroke 16 years earlier, controlled a robotic arm using thoughts conveyed through a brain control interface (BCI) chip to give herself a sip of a chai cinnamon latte for the first time since her injury. Since then, the bandwidth of data obtainable from BCIs has increased exponentially (see figure 8.)

If current BCI trends continue, within the next one to two years we should expect an announcement that a patient with a BCI will be able to type words as fast as the average typist (around 40 words per minute), and within ten years, we should expect world record typing speeds to have been achieved by typing with one's mind (around 200 words per minute.)

FIGURE 8
BRAIN CONTROL INTERFACE (BCI) BANDWIDTH

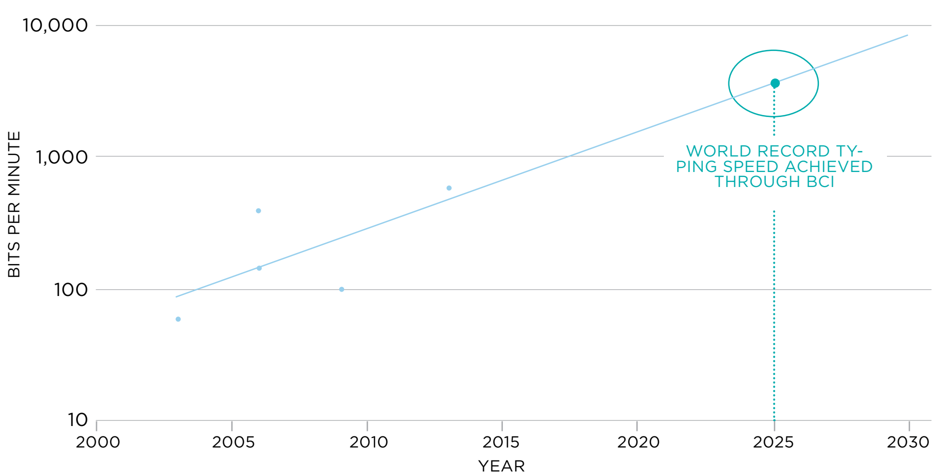

Source: Compiled by the author based on primary sources.

Artificial Melded Intelligence

The exponential progress in BCIs has led some to speculate that the next wave of AI may be artificial melded intelligence (AMI), which combines AI with BCIs. Bryan Johnson, the former CEO of Braintree, who sold his company to eBay for US$700 million, has formed a company called Kernel, which aims to create neuroprosthetics to correct damage caused by strokes or to enhance the human brain. Similarly, Elon Musk has founded a company called Neuralink, which aims to create a BCI linking a human brain with AI which he terms a "neural lace" (Statt, 2016).

THE PATIENT PERSPECTIVE

Going forward, it is expected that progress in the field of healthcare will be driven through progress in AI. Simultaneously, progress in the healthcare sector toward understanding the human mind will drive progress in AI. This synergy can be expected to continue to fuel the exponential trends that have led to the remarkable progress we have seen to date.

Some of the biggest companies today built their success without owning the products they sold or the assets to deliver their services. This has raised the question of what the future of healthcare may look like. The trends in AI shows how the next biggest company in healthcare may not employ any healthcare workers or own healthcare facilities. The company may simply be AI-delivered healthcare from the cloud.

The continuing success of AI in the sector may be good or possibly bad for the those currently employed in it, but early progress suggests it will almost certainly be good for patients.

NOTES

[1]Electronic Freedom Foundation, "Measuring the progress of AI research," https://www.eff.org/ai/metrics.
[2]Howard, J. 2015. Singularity University Exponential Medicine Conference.
[3]Howard, J. 2017. AIMed Conference. Laguna Niguel, California.
[4]See the studies by Sanjay Purushotham atwww-scf.usc.edu/-spurusho/.

REFERENCES

Angier, N. 1990. "Great 15-Year Project to Decipher Genes Stirs Opposition." The New York Times. June 5.
Burgess, M 2016. "Elon Musk's OpenAI and Google's DeepMind Release Their AI Playgrounds to Everyone." Wired. December 5.
Doyle, R. 2016. "Optimising CRISPR Genome Editing using Machine Learning." Re-Work Conference. London.
Funk, L.M., Weiser, T.G., Berry, W.R., et al. 2010. "Global Operating Theatre Distribution and Pulse Oximetry Supply: An Estimation from Reported Data." *The Lancet* 376 (9746): 1055–1061.
Hsu, J. 2017. "AI Learns Gender and Racial Biases from Language." IEEE Spectrum. April 13.
IBM Research. 2012. "IBM Simulates 530 Billion Neurons, 100 Trillion Synapses on Supercomputer." KurzweilAI. November 19.
Kubota, T. 2017. "Deep Learning Algorithm Does as Well as Dermatologists in Identifying Skin Cancer." Stanford News. January 25.
Kurzweil, R. 2014. *How to Create a Mind: The Secret of Human Thought Revealed*. London: Duckworth.
Makary, M.A., and Daniel, M. 2016. "Medical Error—The Third Leading Cause of Death in the US." *BMJ* 353: i2139.
Mannes, J. 2016. "Bryan Johnson Invests $100 Million in Kernel to Unlock the Power of the Human Brain." TechCrunch. October 20.
Metz, C. 2016. "7,500 Faceless Coders Paid in Bitcoin Built a Hedge Fund's Brain." Wired. December 12.
Peng, L., and Gulshan, V. 2016. "Deep Learning for Detection of Diabetic Eye Disease." Google Research Blog. November 29.
Statt, N. 2017. "Elon Musk Launches Neuralink, a Venture to Merge the Human Brain with AI." The Verge. March 27.
The Brain Preservation Foundation. 2017. "Newly Invented Aldehyde-Stabilized Cryopreservation Procedure Wins Brain Preservation Prize." February. Available at: www.brainpreservation.org.
Vedantam, S. 2007. "The Color of Healthcare: Diagnosing Bias in Doctors." The Washington Post. August 13.
Vincent, J. 2017a. "DeepMind's Go-Playing AI Doesn't Need Human Help to Beat Us Anymore." The Verge. October 18.
—. 2017b. "DeepMind's AI Became a Superhuman Chess Player in a Few Hours, Just for Fun." The Verge. December 6.
World Health Organization. 2015."New Report Shows That 400 million Do Not Have Access to Essential Health Services." WHO and World Bank. June 12.
Xiong, W., Droppo, J., Huang, X., et al. 2016. "Achieving Human Parity in Conversational Speech Recognition." Microsoft Research Technical Report No. MSR-TR-2016-71. Microsoft.

CASE STUDY

Precision Medicine

The Watson Robotlution

Juan Hoyos
IBM Latin America

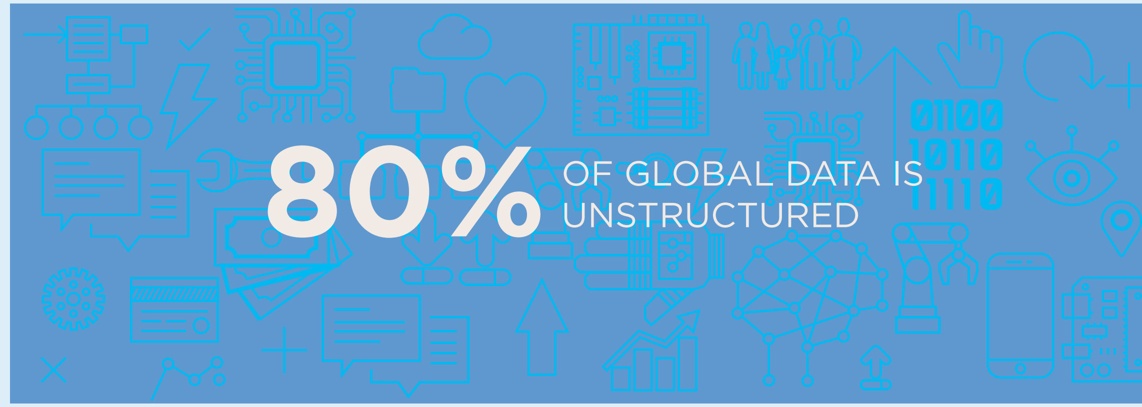

The world has become data-driven. Companies from different industries are transforming themselves digitally while consumers have been empowered and have immediate needs. Innovation now happens in days, not in years. In this context, data explosion has become both the fuel and the result of the digital transformation.

Data has redefined relationships, experiences, habits, business models, and jobs. More than 2.7 zettabytes of data exist in the digital universe today. More than 5 billion people are calling, texting, tweeting, and browsing on mobile phones worldwide. The volume of business data worldwide, across all companies, is estimated to double every 1.2 years (Mulcahy, 2017). Big data is an issue in every business today. The amount of data collected and stored by organizations is increasing exponentially, and the ability to access and analyze it is becoming not only crucial but also a competitive advantage for companies.

While 20% of all the data in the world is searchable through traditional search engines, 80% is not. This unstructured data is comprised of videos, pictures, free text in medical records, social media activity, and so on.

The amount of health-related data is growing exponentially. Densen (2011) show that medical data is expected to double every 73 days by 2020. Each individual is likely to generate more than one million gigabytes of health-related data in a lifetime.

The field of medical scientific literature faces a similar challenge. There are more than 27 million scientific papers on Pubmed. Some 1,261,379 articles were published in 2016. The volume of medical knowledge generated nowadays is outstripping the ability of even the most brilliant clinicians to keep up to date (Corlan, 2016). In 2015, 60 billion medical images were generated in the US (Watson Health, 2016). Radiologists now spend 64% of their time on noninterpretive tasks.

The time needed to analyze and consume this information—locating insights specific to each patient's unique needs to potentially improve treatment outcomes—is more limited than ever. This is especially true when we consider that consultation length in primary care physicians' offices in 67 countries varies between 48 seconds, in Bangladesh, to 22 minutes, in Sweden (Irving et al., 2017).

Through natural language processing and machine learning, artificial Intelligence (AI) and cognitive computing can enable organizations to tap into this 80% of unstructured data and manage it intelligently, transforming it into valuable insights to help professionals make decisions more accurately, based on evidence that is beyond the human brain's capacity to analyze. Some of the most promising use cases for AI and cognitive computing in healthcare include predictive analytics, precision medicine, and clinical decision support.

IBM's Watson Health became a leading name early on in the healthcare industry by using its natural language processing and cognitive computing capabilities to develop clinical decision-support tools that rest on five pillars: oncology and genomics, life sciences, population health, imaging, and the transition to value-based care models.

According to the WHO, cancer is already one of the leading causes of death in the world. As a consequence, oncologists' workload has inevitably been increasing. At the same time, as of September 2015, PubMed contained more than 2.7 million articles filed under the Medical Subject Heading (MeSH) heading "Neoplasms" with more than 100,000 new articles accruing annually since 2011.

To keep up with just the cancer literature, one would have to read 17 articles per hour, 365 days per year (Warner, 2015). It is clear that physicians suffering from information overload need technology-based decision-support systems.

Watson for Oncology is a cognitive decision-support tool that uses natural language processing to ingest patient data in structured and unstructured formats in order to provide ranked, evidence-based treatment options for consideration by the treating oncologists. It was developed by IBM in collaboration with the Memorial Sloan Kettering Cancer Center, one of the most prestigious cancer centers in the world.

After ingesting about 15 million pages of text, including content from textbooks, scientific articles and treatment protocols from the Memorial Sloan Kettering Cancer Center and the National Comprehensive Cancer Network (NCCN) and being trained by oncology experts, the tool is able to rank personalized treatment plans for patients. The tool summarizes key medical attributes of a patient and cross-matches it with peer-reviewed studies from the scientific literature. Then it ranks all the available treatment options for each patient and provides insights about drug administration schemes, as well as supportive care and adverse effects expected for each treatment modality presented. Currently, Watson for Oncology covers 13 cancer types, representing 80% of the global incidence of cancer.

In Somashekhar et al. (2017), a blinded study of 362 cases carried out by the Manipal Comprehensive Cancer Centre in Bangalore, India, Watson for Oncology achieved a concordance rate of 96% for lung, 81% for colon, and 93% for rectal cancer cases compared to recommendations from the multidisciplinary tumor board at the same center.

Somashekhar et al. (2018), another study led by oncologists at Manipal Hospitals in India, Watson was found to be concordant with the oncologists' own treatment decisions in 93% of 638 breast cancer cases. Watson for Oncology has also achieved a concordance rate of 83% for multiple cancer types compared to recommendations from oncologists in Suwanvecho et al. (2017), a study carried out at Bumrungrad International Hospital, a multispecialty hospital in Bangkok, Thailand.

In a qualitative study, oncologists in Mexico (Sarre-Lazcano et al., 2017) found Watson for Oncology to be useful to help them identify potential treatment options for their patients, particularly in clinics that lack subspecialist expertise, and for training medical students and residents.

Genomic sequencing in cancer care is a promising and recurrent topic in discussions regarding precision medicine. Although the cost of genome sequencing has dropped significantly in the latest years, it is still challenging to analyze the amount of data it generates. Multidisciplinary decision-support teams still need to manually analyze gene panels to determine the effect of mutations and make recommendations so that treating oncologists can prioritize clinical trial options for cancer therapy that is personalized to the patient and his or her unique genome (Khotskaya, Mills, and Mills Shaw, 2017).

Watson for Genomics is a cloud-based solution that helps tackle this challenge through the power of cognitive computing. It empowers oncologists to deliver precision medicine to their patients by providing evidence-based genetic tumor sequencing analysis and compiling possible therapeutic options. Watson for Genomics uses relevant information extracted from the massive volumes of medical literature to provide a genetic analysis of a patient's treatable cancer-causing mutations. The report generated by Watson includes recommendations for potential targeted therapies that are relevant to the unique DNA profile of a patient's tumor. Clinicians can then evaluate the evidence to determine whether targeted therapy may be more effective than other options.

In a recent comparison study at the New York Genome Center, Hwang et al. (2017) used Watson for Genomics to help scale the interpretation of whole genome sequencing They found that the tool could provide a report of potentially actionable genomic insights for one patient in just 10 minutes, as compared to the 160 hours needed to arrive at similar conclusions through manual interpretation.

In a retrospective analysis (Patel et al., 2017) of 1,018 cancer cases at the UNC Lineberger Cancer Center, the molecular tumor board identified actionable genetic alterations in 703 cases, which Watson also confirmed. However, Watson for Genomics was able to identify additional potential therapeutic options in 323 patients, or one-third of the cases reviewed, which the molecular tumor board had not identified. Of these, 96 had not been previously identified as having an actionable mutation.

Cognitive computing and AI tools are becoming fundamental for the empowerment of physicians in the face of the data overload that exists in healthcare. As ironic as it may seem, in a world where the amount of data to be considered for proper care of a patient exceeds the processing capacity of the human brain, technology is beginning to emerge as an important, perhaps fundamental, ally for helping healthcare professionals deliver more humane and personalized care to patients.

REFERENCES

Corlan, A.D. 2016. "Medline Trend: Automated Yearly Statistics of PubMed Results for any Query, 2004." Last changed: February 14, 2012. http:/dan.corlan.net/medline-trend.html.

Densen, P. 2011. "Challenges and Opportunities Facing Medical Education." *Transactions of the American Clinical and Climatological Association* 122: 4858.

Hwang, Y.T., Lakshmanan, R., Davagnanam, I., et al. 2017. "Brainstem Phenotype of Cathepsin A-Related Arteriopathy with Strokes and Leukoencephalopathy." *Neurology Genetics* 3 (4) e165.

Irving, G., Neves, A.L., Dambha-Miller, H., et al. 2017. "International Variations in Primary Care Physician Consultation Time: A Systematic Review of 67 Countries." *BMJ Open* 7: e017902.

Khotskaya, Y.B., Mills, G.B., and Mills Shaw, K. R. 2017. "Next-Generation Sequencing and Result Interpretation in Clinical Oncology: Challenges of Personalized Cancer Therapy." *Annual Review of Medicine* 14 (68): 113–125.

Mulcahy, M. 2017. "Big Data Statistics & Facts for 2017." Waterford Technologies Blog. February 22.

Patel, N.M., Michelini, V.V., Snell, J.M., et al. 2017. "Enhancing Next-Generation Sequencing-Guided Cancer Care Through Cognitive Computing." *The Oncologist* 22: 1–7.

Sarre-Lazcano, C., Alonso, A.A., Melendez, F.D.H., et al. 2017. "Cognitive Computing in Oncology: A Qualitative Assessment of IBM Watson for Oncology in Mexico." *Journal of Clinical Oncology* 35 (15).

Somashekhar, S.P., Sepúlveda, M.J., Norden, A.D., et al. 2017. "Early Experience with IBM Watson for Oncology (WFO) Cognitive Computing System for Lung and Colorectal Cancer Treatment." *Journal of Clinical Oncology* 35 (15): 8527

Somashekhar, S.P., Sepúlveda, M.J., Puglielli, S., et al. 2018. "Watson for Oncology and Breast Cancer Treatment Recommendations: Agreement with an Expert Multidisciplinary Tumor Board." *Annals of Oncology* 29 (2): 418–423.

Suwanvecho, S., Suwanrusme, H., Sangtian, M. et al. 2017. "Concordance Assessment of a Cognitive Computing System in Thailand." *Journal of Clinical Oncology* 35 (15): 6589.

Warner, J.L. 2015. "Grappling with the Data Explosion in Oncology." *Oncology & Hematology Review* 11 (2): 102–103.

Watson Health. 2016. Imagine Your World with Watson. Cambridge: IBM.

ALL ABOUT THE FUTURE OF REGIONAL INTEGRATION

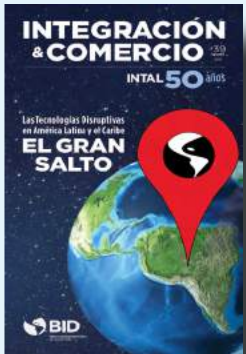

ACCELERATED DEVELOPMENT
A detailed analysis of the impact of exponential new technologies on Latin American economies

MARKET DIVERSIFICATION
Concrete ideas for strengthening trade and cooperation between India and Latin America

SOUTH-SOUTH INTEGRATION
The new normal for the Chinese economy and the future of the country's relations with Latin America and the Caribbean

INDUSTRY 4.0
Digital transformation in different productive sectors: regulatory challenges and industrial policies

SUSTAINABLE GROWTH
Inspired by the papal encyclical *Laudato Si'*, experts from around the world put forward concrete ideas for moving toward an ecofriendly form of integration

WHAT DOES THE NEXT GENERATION THINK?
How young Argentinians relate to technology, the digital economy, and the labor market of the future

CUTTING-EDGE KNOWLEDGE

Multidisciplinary experts from around the world analyze the impact of artificial intelligence on Latin America and the Caribbean from a strategic, holistic perspective.

• PETER DIAMOND
Nobel Prize in Economic Sciences, 2010

• DAVE DONALDSON
John Bates Clark Medal for the best US economist under the age of 40

• GEOFF MULGAN
CEO of the United Kingdom's National Endowment for Science, Technology and the Arts (NESTA)

• RAFAEL REIF
President of the Massachusetts Institute of Technology (MIT)

• NICOLAS MIAILHE
Cofounder and president of The Future Society at Harvard University

• KARIM LAKHANI
Professor of Business Administration at Harvard Business School

• LUIS ASCENCIO
International consultant on digital ports for the Latin American and Caribbean Economic System (SELA)

• AVINASH VASHISTHA AND ANKITA VASHISTHA
Founder and researcher at Tholons Inc., a global firm focusing on digital transformation and service exports

• JOHN ATKINSON
Director of the IT Center at Adolfo Ibañez University, Chile

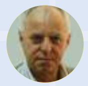

• DANIEL HEYMANN AND PABLO MIRA
Professors and researchers at the University of Buenos Aires. Experts in macroeconomics and behavioral economics

• YOLANDA LANNQUIST
Researcher at the Harvard Kennedy School's AI Initiative

• SANTIAGO CHELALA
INTAL/IDB consultant and author of the Synthetic Indicator for the Risk of Job Automation

• ANAND RAO
Global artificial intelligence lead at PriceWaterhouseCoopers

• MICHAEL GILLAM
CEO of HealthLab and professor at Singularity University, where he leads initiatives on the use of virtual reality and robotics in medicine

• WELBER BARRAL
Former secretary for foreign trade of Brazil

• PHIL TINN AND MICHAEL LIN
Researchers at MIT LAB and experts on new forms of transportation

• JULIÁN SIRI
Professor at the University of CEMA, Argentina, and holder of a master's in Financial Engineering from Columbia University

• JUAN LAVISTA
Senior director of data science at Microsoft and editor of the Microsoft Journal of Applied Research (MSJAR)

• GABRIEL PETRUS
Executive director of ICC Brazil and managing director of the Intelligent Tech & Trade Initiative for the use of AI in trade negotiations

• DÉBORA SCHAPIRA
Professor at the National University of Tres de Febrero, Argentina. Expert on education policies and AI applications for education

• EDUARDO PLASTINO
Senior researcher at Accenture Research

• MARCOS HERRERA
Specialist in spatial econometrics, CONICET, Argentina

• ROSA GONZÁLEZ-RAMÍREZ
Professor at the School of Engineering, University of the Andes, Chile, and a specialist in port logistics

• NATALIA SAMPIETRO
Director of public data at Argentina's Ministry of Modernization

• JUAN CORVALÁN
Assistant federal prosecutor for the City of Buenos Aires. Promoted the implementation of the AI application Prometea within the city's legal system

• CARLOS CHESÑEVAR AND ELSA ESTEVEZ
Professors at the National University of the South, Argentina. Leading Latin American experts on artificial intelligence

• JULIA NUÑEZ TABALES, JOSÉ CARIDAD Y OCERÍN AND MARIA GARCÍA-MORENO
Researchers at the University of Córdoba, Spain

• ANTON KORINEK
Professor at Johns Hopkins University and an expert on innovation and income distribution

• SAMIR SARAN AND MADHULIKA SRIKUMAR
President of and expert on cybersecurity at the Observer Research Foundation, India

• DOAA ABU ELYOUNES
Expert on justice and AI, Harvard University

• RICARDO RODRÍGUEZ
Scientist at the School of Exact Science, University of Buenos Aires

• KATE POUNDER
CEO of Alphabeta Australia, responsible for the strategic report on job conversion and the digital economy

• JUAN HOYOS
Latin America executive leader, IBM Watson

INTAL-LAB

A MEETING PLACE FOR THE COCREATION OF INNOVATIVE IDEAS ON INTEGRATION AND TRADE IN LATIN AMERICA AND THE CARIBBEAN

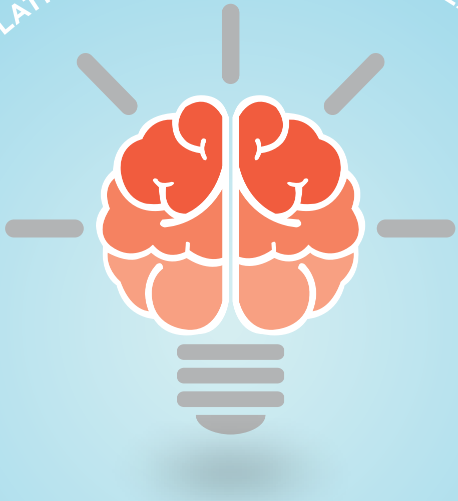

THE INSTITUTE FOR THE INTEGRATION OF LATIN AMERICA AND THE CARIBBEAN HAS STARTED INTAL-LAB, A SPACE WHERE WE CAN REIMAGINE THE REGION'S FUTURE TOGETHER. PRESENTATIONS FROM GLOBAL EXPERTS, HACKATHONS, INNOVATION RALLIES, CROWDWORKING EXPERIMENTS, SEMINARS, AND DESIGN THINKING ARE SOME OF THE MANY WAYS THAT INTAL IS SEEKING TO FOSTER COLLECTIVE INTELLIGENCE THAT IMPROVES LIVES THROUGH INTEGRATION AND TRADE.

IDB

WWW.IADB.ORG/INTAL

INTEGRATION & TRADE

#44
July
2018

Whether we know it or not, we are all living on a new planet: Planet Algorithm. This is a cyberphysical space in which billions of pieces of data are transported at hyper-speed and are analyzed by increasingly sophisticated artificial intelligence (AI) systems. These use algorithms to generate learning and self-learning processes that are making an exponential impact on industry, trade, services, and multiple aspects of our lives together.

In this INTAL/IDB report, over 40 high-profile international experts analyze the risks and opportunities come with the use of intelligent machines in areas that have serious implications for Latin America's productive profile and global role. These range from the possibility of predicting trade negotiation outcomes, commodity prices, and consumer trends to the development of algorithms for use in factories, personalized medicine, extended education, infrastructure prototyping, autonomous ecotransportation, precision agriculture, energy consumption, the legal system, and macroeconomic analysis. They also explore the ethical and equality-related challenges these transformations are posing.

We are witnessing the rise of a technology that is becoming a new factor of production. Artificial intelligence, if guided by a thoughtful, up-to-date, humanist vision, could contribute to consolidating a predictive and inclusive form of regional integration that benefits all Latin Americans.

IDB INTAL