



**Stanford – Vienna
Transatlantic Technology Law Forum**

A joint initiative of
Stanford Law School and the University of Vienna School of Law



TTLF Working Papers

No. 63

**Reinforcement Learning Patents: A
Transatlantic Review**

Brian Haney

2020

TTLF Working Papers

Editors: Siegfried Fina, Mark Lemley, and Roland Vogl

About the TTLF Working Papers

TTLF's Working Paper Series presents original research on technology, and business-related law and policy issues of the European Union and the US. The objective of TTLF's Working Paper Series is to share "work in progress". The authors of the papers are solely responsible for the content of their contributions and may use the citation standards of their home country. The TTLF Working Papers can be found at <http://tflf.stanford.edu>. Please also visit this website to learn more about TTLF's mission and activities.

If you should have any questions regarding the TTLF's Working Paper Series, please contact Vienna Law Professor Siegfried Fina, Stanford Law Professor Mark Lemley or Stanford LST Executive Director Roland Vogl at the

Stanford-Vienna Transatlantic Technology Law Forum
<http://tflf.stanford.edu>

Stanford Law School
Crown Quadrangle
559 Nathan Abbott Way
Stanford, CA 94305-8610

University of Vienna School of Law
Department of Business Law
Schottenbastei 10-16
1010 Vienna, Austria

About the Author

Brian Haney is a writer, entrepreneur, and technology law researcher from Pittsburgh, Pennsylvania. Brian's research focus is empirical intellectual property analysis for high technologies including machine learning, biotechnology, and quantum computing. More specifically, Brian writes and researches about the law as it relates to technical convergence at the cutting edge in biological, synthetic, and artificial intelligence. Brian is a graduate of Notre Dame Law School and formerly practiced as a lawyer in Pittsburgh, Pennsylvania, as well as before the Court of Federal Claims in the District of Columbia. His work has been published in several technology law and intellectual property scholarly journals. Brian has been a TTLF Fellow since 2020.

General Note about the Content

The opinions expressed in this paper are those of the author and not necessarily those of the Transatlantic Technology Law Forum or any of its partner institutions, or the sponsors of this research project.

Suggested Citation

This TTLF Working Paper should be cited as:

Brian Haney, Reinforcement Learning Patents: A Transatlantic Review, Stanford-Vienna TTLF Working Paper No. 63, <http://tflf.stanford.edu>.

Copyright

© 2020 Brian Haney

Abstract

One of the most difficult problems with developing scalable artificial intelligence (AI) is tracking technical innovation. As such, this research will provide empirical data relating to patents with legal claims to state of the art in AI technologies, reinforcement learning. A keystone architecture in the machine learning paradigm, reinforcement learning technologies power trading algorithms, driverless cars, and space satellites.

In competitive global markets, owning and protecting legal rights to reinforcement learning technologies may mean the difference between market dominance and irrelevance for the transatlantic firm. Aggregating patent data for reinforcement learning technologies from both the United States Patent and Trademark Office (USPTO) and the European Patent Office (EPO), this research offers a transatlantic perspective on reinforcement learning innovation and intellectual property protections.

Acknowledgements

Thanks to Angela Elias, Brad Haney, Broderick Haney, Mark Lemley, Roland Vogl, Siegfried Fina, Fernando Martinez, Branden Keck, Justin Goodwill and Brian Bozzo.

Table of Contents

<i>Introduction</i>	2
<i>I. Reinforcement Learning Software</i>	3
A. Statistical Simulations	4
B. Deep Decisions	10
C. Applied Art	14
<i>II. Patent Data</i>	16
A. USPTO	17
B. EPO	18
C. Transatlantic	18
<i>III. Comparative Claims Critique</i>	19
A. Technically Translating	20
B. Simplifying Syntax	23
C. Scaling Scope	28
<i>IV. Property Strategy</i>	33
A. Patent Protections	33
B. Software Secrets	38
C. Code Copyrights	41
<i>Conclusion</i>	46
<i>Appendices</i>	48
Appendix A. Notation Summary	48
Appendix B. Global Statistics	49
Appendix C. Top Ten Most Valuable Patents	50
Appendix D. Transatlantic Translations	51

Introduction

Reinforcement learning is a machine learning¹ technique for computer programs to achieve goals. For example, reinforcement learning algorithms drive cars, land rockets, and diagnose patients. As a field of research, reinforcement learning rests at the intersection of computer science, philosophy, and mathematics.² A vital component for global AI development, reinforcement learning software is the keystone to artificial general intelligence.³

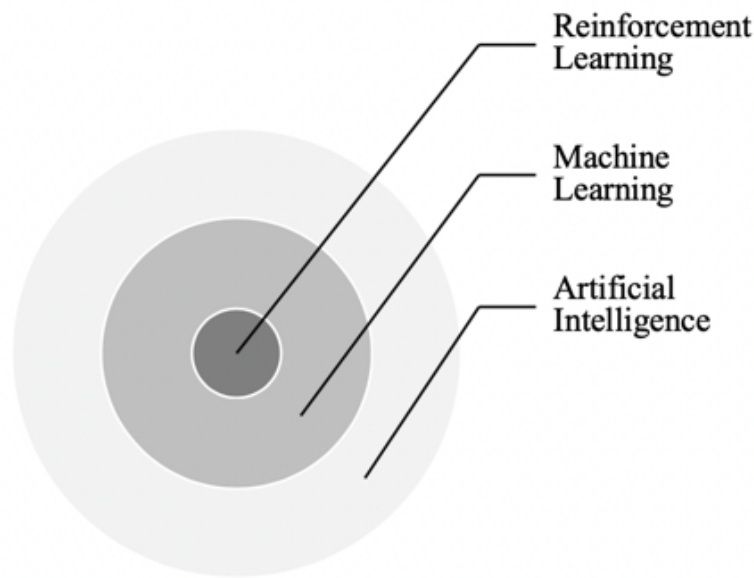


Figure 1

Protecting the property rights to reinforcement learning technologies is the differentiating factor for firms advancing AI technology. Despite some more aggressive strategies with a heavy

¹ Emily Berman, *A Government of Laws and Not of Machines*, 98 B.U. L. REV. 1277, 1278 (2018), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3098995. (Machine learning is a strand of artificial intelligence that sits at the intersection of computer science, statistics, and mathematics, and it is changing the world.) See also Mark A. Lemley, Bryan Casey, *Fair Learning*, 3 (2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3528447. (“The vast potential of ML systems is matched only by their appetite for data.”)

² Leslie Pack Kaelbling, et al., *Reinforcement Learning: A Survey*, J. of Artificial Intelligence Research (1996), <http://www.cse.msu.edu/~cse841/papers/kaelbling.pdf>. (Surveying the field of reinforcement learning.)

³ Maria Schuld, et al., *An introduction to quantum machine learning* 4 (2014), <https://arxiv.org/pdf/1409.3097.pdf>. (“Reinforcement learning is a central mechanism in the development and study of intelligent agents.”)

emphasis on software patents, the majority deploy a mixed intellectual property strategy.⁴ The word patent comes from the Latin word *paten* – meaning visible.⁵ And, some firms have taken an open strategy in reinforcement learning software development.⁶

This Article proceeds in four parts. Part I defines, discusses and details reinforcement learning software architectures and design considerations. Part II introduces a transatlantic patent dataset for reinforcement learning technologies. Part III compares and contrasts patent claims from the United States and Europe. And, Part IV discusses intellectual property strategy and the protections available for reinforcement learning technologies including patents, trade secrets and copyrights. In sum, this Article contributes the first empirical transatlantic patent review for reinforcement learning technologies.

I. Reinforcement Learning Software

Reinforcement learning software optimizes agent performance according to a reward.⁷ The process involves building models and developing systems for decision making embedded in software programs.⁸ Reinforcement learning algorithms contain three elements: (1) model: the description of the agent-environment relationship;⁹ (2) reward: the agent’s goal;¹⁰ and (3) policy:

⁴ Stefania Fusco, et al., Monetization Strategies of University Patents Through PAEs: an Analysis of US Patent Transfers, ISSI Conference Proceedings (2019), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3410086.

⁵ Oxford Latin Desk Dictionary.

⁶ Jeanne C. Fromer, *Machines as the New Oompa-Loompas: Trade Secrecy, the Cloud, Machine Learning, and Automation*, N.Y.U. L.R., 706, 717 (2019), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3359746.

(“Contractual restrictions on disclosure and reverse engineering diminish an otherwise critical intrinsic weakness of trade secrecy for software: software’s openness to independent discovery and reverse engineerability.”)

⁷ MYKEL J. KOCHENDERFER, DECISION MAKING UNDER UNCERTAINTY 77 (2015). *See also* United States Patent No. 10,346,741 to Mnih, et al. Asynchronous deep reinforcement learning (July 9, 2019)(Assigned to DeepMind Technologies – a Google subsidiary.)

⁸ Leslie Pack Kaelbling, *Learning in Embedded Systems* (1990), <https://apps.dtic.mil/dtic/tr/fulltext/u2/a323936.pdf>.

⁹ Katerina Fragkiadaki, *Deep Q Learning*, Carnegie Mellon Computer Science, CMU 10703 (Fall 2018), https://www.cs.cmu.edu/~katf/DeepRLFall2018/lecture_DQL_katf2018.pdf.

¹⁰ MAXIM LAPAN, DEEP REINFORCEMENT LEARNING HANDS-ON 3 (2018).

the decision function.¹¹ In practice, engineering reinforcement learning systems is a meticulous, time consuming, and data-intensive task process. But the effort is worthwhile because reinforcement learning learns without supervision.¹² In short, the goal of reinforcement learning is to statistically select the policy which maximizes expected reward for an agent acting in an environment.¹³

A. Statistical Simulations

Reinforcement learning's inception may be traced to the year 1846 with the late mathematician, P.L. Chebyshev.¹⁴ Chebyshev's work was a heavy influence on Andrei Markov's work in probability theory, which resulted in one of the world's most important ideas, the Markov Decision Process (MDP).¹⁵ In short, the MDP is a statistical tool for predicting the future. MDPs trace the probabilistic transitions from one state to another through time.¹⁶

¹¹ U.S. Patent No. 9,298,172 Method and apparatus for improved reward-based learning using adaptive distance metrics, Tesauro, et al. (March 29, 2016) (Assigned to International Business Machines Corporation). See also Katerina Fragkiadaki, Deep Q Learning, Carnegie Mellon Computer Science, CMU 10703 (Fall 2018), https://www.cs.cmu.edu/~katef/DeepRLFall2018/lecture_DQL_katef2018.pdf.

¹² JOHN D. KELLEHER, DEEP LEARNING 28-29 (2018).

¹³ Barry, Jennifer, Daniel T. Barry, and Scott Aaronson, *Quantum Partially Observable Markov Decision Processes*, American Physical Society, MIT Open Access Articles at 2 (2014), <https://journals.aps.org/pr/abstract/10.1103/PhysRevA.90.032311>.

¹⁴ P.L. Chebyshev, *Démonstration élémentaire d'une proposition Générale de la théorie des probabilités*, 33 J. Reine Angew. Math. 259 (1846). See also Paul Butzer, *P. L. Chebyshev (1821-1894) A Guide to his Life and Work*, 96 Journal of Approximation Theory 111, 118-119 (1999). ("Remaining unnoticed at the time, this paper had thus no effect on the controversy about laws of large numbers then going on in France.")

¹⁵ Markov was a prominent figure in his time, but his greatest influence was delayed over a century. Claude Shannon was the first to MDP variants in English. Shannon used the MDP to model statistical communication and natural language generation. Today, Markovian techniques pervade the science of modern information theory. Markov's models are used in search algorithms, machine translation, and financial trading. And, the *Markov Decision Process* (MDP) remains the foundation of reinforcement learning. See Gely P. Basharin, et. al, *The Life and Work of A.A. Markov*, 386 Linear Algebra and its Applications 1, 15 (2004). See also GEORGE GILDER, LIFE AFTER GOOGLE 75 (2018). See C.E. Shannon, *A Mathematical Theory of Communication*, Bell Systems Technical Journal 8 (1948). (Referring the stochastic processes of a Markoff process.) See also M. Fréchet, *Méthode des fonctions arbitraires. Théorie des événements en chaîne dans le cas d'un nombre fini d'états possibles*. Paris, Gauthier-Villars (1938).

¹⁶ Until Markov, probability theory avoided temporal considerations. Markov's brilliance was realized in his ability to describe the temporal dependencies between events across time. See MYKEL J. KOCHENDERFER, DECISION MAKING UNDER UNCERTAINTY 77 (2015). See also U.S. Patent No 9,858,171, to Chu, Application analytics reporting (January 2, 2018)(Assigned to Google).(Discussing the use of Markov chains for state transitions.) See also Albert Einstein, *Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen*, 322 Annalen der Physik 549 (1905), <https://www.zbp.univie.ac.at/dokumente/einstein2.pdf>. See also Jeanne C. Fromer, *Machines as the New Oompa-*

Formally, reinforcement learning is described through an agent-environment interaction, with the MDP.¹⁷ Figure 2 describes the agent-environment interaction in an MDP.



Figure 2¹⁸

In an MDP, the interaction begins when an agent chooses an action in the environment's initial state.¹⁹ The model continues to the next state, where the agent receives a reward and a set of actions from which to choose, the agent selects an action, the environment returns a reward and the next state.²⁰ This process continues perpetually until the environment's final state.²¹

Ultimately, in reinforcement learning an agent learns to take actions optimizing a reward.²²

Loompas: Trade Secrecy, the Cloud, Machine Learning, and Automation, N.Y.U. L.R., 706, 720 (2019), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3359746. ("The critical ingredients of machine learning are relevant data and statistical techniques.")

¹⁷ Fabian Ruehle, *Data Science Applications to String Theory*, Physics Reports at 134 (2019), <https://doi.org/10.1016/j.physrep.2019.09.005>. See also Serena Yeung, et al., End-to-end Learning of Action Detection from Frame Glimpses in Videos (2016), https://www.cv-foundation.org/openaccess/content_cvpr_2016/html/Yeung_End-To-End_Learning_of_CVPR_2016_paper.html. (Introducing a fully end-to-end approach for action detection in videos that learns to directly predict the temporal bounds of actions.)

¹⁸ Leslie Pack Kaelbling, et al., Planning and acting in partially observable stochastic domains, 101 Artificial Intelligence 99 (1998), <https://people.csail.mit.edu/lpk/papers/aij98-pomdp.pdf>.

¹⁹ EUGENE CHARNAK, INTRODUCTION TO DEEP LEARNING 113 (2018).

²⁰ See Katerina Fragkiadaki, et al., Grouping-Based Low-Rank Trajectory Completion and 3D Reconstruction (2014), <https://www.cs.cmu.edu/~katf/nrsfm.html>. (Exploring how an agent can be equipped with an internal model of the dynamics of the external world, and how it can use this model to plan novel actions by running multiple internal simulations.)

²¹ Volodymyr Mnih et al., Human-Level Control Through Deep Reinforcement Learning, 518 NATURE INT'L J. SCI. 529, 529 (2015).

²² Barry, Jennifer, Daniel T. Barry, and Scott Aaronson, *Quantum Partially Observable Markov Decision Processes*, American Physical Society, MIT Open Access Articles at 2 (2014), <https://journals.aps.org/pr/abstract/10.1103/PhysRevA.90.032311>. See also U.S. Patent No. 10423129 to Huynh, et al. Controlling dynamical systems with bounded probability failure (September 24, 2019). (Assigned to Massachusetts Institute of Technology)

For reinforcement learning software, the environment²³ defines the problem.²⁴ In other words, dynamic environments are made of continuous or discrete states, representing moments in time.²⁵ For example, in robotics control systems, the environment is made up of states for moments in time in which the environment exists. Alternatively, in a trading algorithm the environment may be made up of a portfolio of stocks.²⁶

An agent is an algorithm solving the environment or problem.²⁷ The goal for any agent in an MDP is to maximize expected reward during the episode.²⁸ In other words, the agent's goal is to maximize its total reward, rather than the reward for its immediate state.²⁹ The reward is a method of teaching the agent what it should do and is meant to formalize the idea of a goal.³⁰ The goal would be to allow the agent to make sacrifices for a particular move, reducing

²³ Environments are made up of two types of space, state spaces and action spaces. There are two types of state spaces. The first is fully observable. In a fully observable environment, the agent directly senses the total state at each time step in the environment. By contrast, in a partially observable environment, the agent senses a fraction of the environment. Despite the environment's partial observability, performance must not suffer, and the agent's actions must be ensuring passenger safety. The total of all the states in an environment is called the episode, which concludes with the last state, the terminal state. *See* Barry, Jennifer, Daniel T. Barry, and Scott Aaronson, *Quantum Partially Observable Markov Decision Processes*, American Physical Society, MIT Open Access Articles at 2 (2014), <https://journals.aps.org/pr/abstract/10.1103/PhysRevA.90.032311>. *See also* U.S. Patent No. 8,060,454 to Das, et al., Method and apparatus for improved reward-based learning using nonlinear dimensionality reduction (November 15, 2011). (Assigned International Business Machines Corporation)

²⁴ MAXIM LAPAN, DEEP REINFORCEMENT LEARNING HANDS-ON 8 (2018). *See also* U.S. Patent No. 9,298,172 Method and apparatus for improved reward-based learning using adaptive distance metrics, Tesauro, et al. (March 29, 2016) (Assigned to International Business Machines Corporation).

²⁵ MAXIM LAPAN, DEEP REINFORCEMENT LEARNING HANDS-ON 20 (2018).

²⁶ MAXIM LAPAN, DEEP REINFORCEMENT LEARNING HANDS-ON 217 (2018). *See also* Volodymyr Mnih, Koray Kavukcuoglu, Methods and Apparatus for Reinforcement Learning, U.S. Patent No. 9,679,258 B2 (2017) (<https://patents.google.com/patent/US9679258B2/en>).

²⁷ The agent may iterate over the action space, selecting actions according to a defined policy. *See* EUGENE CHARNAK, INTRODUCTION TO DEEP LEARNING 113 (2018). *See also* United States Patent No. 10,498,855 to Latapie, et al., Contextual services in a network using a deep learning agent (December 3, 2019). (Assigned to Cisco Technology, Inc.)

²⁸ Episode refers to the total experience of an agent progressing through an environment a terminal state. *See* MYKEL J. KOCHENDERFER, DECISION MAKING UNDER UNCERTAINTY 77 (2015). *See also* United States Patent No. 10,498,855 to Latapie, et al., Contextual services in a network using a deep learning agent (December 3, 2019). (Assigned to Cisco Technology, Inc.)

²⁹ EUGENE CHARNAK, INTRODUCTION TO DEEP LEARNING 113 (2018).

³⁰ EUGENE CHARNAK, INTRODUCTION TO DEEP LEARNING 113 (2018).

immediate reward, at the expense of increasing the probability of winning the overall game, the total reward.

Defining the reward for a reinforcement learning system is often one of the most challenging aspects of algorithmic development.³¹ The reward is easier to describe for a task like missile control, where the agent need only take actions to minimize the missile's distance from the target.³² An important distinction in reinforcement learning is the relationship between reward and value.³³ The reward defines the response from taking an action in a given state, where the value refers to the total amount of reward over an episode.³⁴ The rewards are used to update the agent's knowledge over time, so it learns to take actions returning the highest rewards.³⁵ For each time step, the reward is a number associated with a corresponding action.³⁶ In other words, reward is a measure of short-term gain and value is a measure of long-term reward.³⁷ The agent's policy determines the value the agent returns over the course of an episode.³⁸

³¹ NICK BOSTROM, SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES 239 (2017). *See also* U.S. Patent No. 10,467,274, to Ren, et al. Deep reinforcement learning-based captioning with embedding reward (November 5, 2019). (Assigned to Snap Inc.)

³² Rebecca Crootof, *Autonomous Weapons Systems and the Limits of Analogy*, 9 HARV. NAT'L SEC. J. 51, 59 (2018). *See also* Shixun You, et al., *Deep Reinforcement Learning for Target Searching in Cognitive Electronic Warfare*, IEEE Access Vol. 7, 37432, 37438 (2019).

³³ *See Generally* Serena Yeung, et al., Learning to Learn from Noisy Web Videos (2017), http://openaccess.thecvf.com/content_cvpr_2017/html/Yeung_Learning_to_Learn_CVPR_2017_paper.html. (Using Q-learning to learn a data labeling policy on a small labeled training dataset, and then uses this to automatically label noisy web data for new visual concepts.)

³⁴ EUGENE CHARNIAK, INTRODUCTION TO DEEP LEARNING 113-114 (2018).

³⁵ MYKEL J. KOCHENDERFER, DECISION MAKING UNDER UNCERTAINTY 77 (2015).

³⁶ *See* U.S. Patent No. 8,429,096 to Soundararajan, et al. Resource isolation through reinforcement learning (April 23, 2013). (Assigned to Amazon Technologies, Inc.)

³⁷ EUGENE CHARNIAK, INTRODUCTION TO DEEP LEARNING 113-114 (2018).

³⁸ Volodymyr Mnih, Koray Kavukcuoglu, Methods and Apparatus for Reinforcement Learning, U.S. Patent Application No. 14/097,862 at 5 (filed Dec. 5, 2013), <https://patents.google.com/patent/US20150100530A1/en>.

A policy³⁹ is a mapping from states to probabilities for selecting actions.⁴⁰ In other words, a policy is the way in which an agent makes decisions.⁴¹ The goal for reinforcement learning is to develop a policy allowing the agent to maximize the value it returns for a given episode.⁴² One of the main challenges in reinforcement learning is balancing exploring for new rewards and exploiting learned rewards.⁴³ In other words, an agent must prefer actions it has found to be effective in producing rewards, but it also must try new actions to discover the environment's best rewards.⁴⁴ So, the agent has to exploit its knowledge to gain rewards, but also has to explore to take better actions in the future.⁴⁵ Thus, the agent tries a variety of actions, both stochastically and deterministically, progressively favoring those that return the best value.⁴⁶

³⁹ Formally, the policy is represented as π . In general, there are two types of policies, deterministic and stochastic policies. In a deterministic policy, the state determines the action

$$a = \pi(s)$$

In a stochastic policy, the agent randomly decides each action:

$$\pi(a|s) = \mathbb{P}[a|s].$$

The goal for a given environment is to find the optimal policy. *See also* U.S. Patent No. 8,478,642, System, method and device for predicting navigational decision-making behavior (July 2, 2013). (Assigned to Carnegie Mellon University). *See also* U.S. Patent No. 10,146,286, to Lee, et al., Dynamically updating a power management policy of a processor (December 4, 2018). (Assigned to Intel Corporation.)

⁴⁰ MYKEL J. KOCHENDERFER, DECISION MAKING UNDER UNCERTAINTY 80 (2015). *See also* U.S. Patent No. 8,060,454 to Das, et al., Method and apparatus for improved reward-based learning using nonlinear dimensionality reduction (November 15, 2011). (Assigned International Business Machines Corporation)

⁴¹ MYKEL J. KOCHENDERFER, DECISION MAKING UNDER UNCERTAINTY 80 (2015).

⁴² Jeanne C. Fromer, Learning Optimal Discourse Strategies in a Spoken Dialogue System, MIT Masters Thesis, 40 (1998), <https://dspace.mit.edu/bitstream/handle/1721.1/47703/42306186-MIT.pdf?sequence=2&isAllowed=y>. (“These algorithms can calculate optimal discourse policies for Markov decision problems (MDPs), accessible, stochastic environments with a known transition model.”)

⁴³ Fang Liu, Assessment of Bayesian Expected Power via Bayesian Bootstrap 14 (2017), <https://arxiv.org/abs/1705.04366>. (“The bootstrap-based procedures will appeal to non-Bayesian practitioners given their analytical and computational simplicity and easiness in implementation.”)

⁴⁴ MARVIN MINSKY, SOCIETY OF MIND, 76 (198).

⁴⁵ U.S. Patent No. 7,395,252 to Anderson, et al., Innervated stochastic controller for real time business decision-making support (July 1, 2008) (Assigned to The Trustees of Columbia University in the City of New York).

⁴⁶ U.S. Patent No. 8,468,041 to Vengerov, Using reinforcement learning to facilitate dynamic resource allocation (June 18, 2013) (Assigned to Oracle America, Inc.) *See also* U.S. Patent No. 10,296,004 to Nishin, Autonomous operation for an autonomous vehicle objective in a multi-vehicle environment (May 21, 2019). (Assigned to Toyota)

Generally, an optimal policy is developed to maximize value.⁴⁷ A value function⁴⁸ is used to compute the value of a given state according to a defined policy.⁴⁹ Policy evaluation is the process of computing the expected reward from executing a policy in a given environment.⁵⁰ Policy evaluation can be used in a general process called policy iteration⁵¹ for computing an optimal policy.⁵² Policy iteration is effective because the number of policies for an agent in an MDP are finite.⁵³ Thus, the iterative process of updating policies must converge to an optimal policy and optimal value function in a finite number of iterations.⁵⁴

⁴⁷ PAUL JOHN WERBOS, THE ROOTS OF BACKPROPAGATION FROM ORDERED DERIVATIVES TO NEURAL NETWORKS AND POLITICAL FORECASTING 306 (1994).

⁴⁸ A value function computes state value according to a policy. For example, the value function $V^\pi(s)$ is equal to the expected sum of the discounted rewards for executing policy π :

$$V^\pi(s) = E_{T \sim \pi} \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1}) | s_0 = s \right]$$

The expected future rewards are discounted with a discount factor γ . The discount factor is typically defined:

$$0 < \gamma < 1.$$

If the discount factor is low, the agent considers present rewards to be worth more and if the discount factor is high, future rewards are worth more – relatively speaking. *See* Ahmad El Sallab et al., Deep Reinforcement Learning Framework for Autonomous Driving, (Apr. 8, 2017), (unpublished paper) (accessed at <https://arxiv.org/pdf/1704.02532.pdf>). *See also* U.S. Patent No. 8,060,454 to Das, et al., Method and apparatus for improved reward-based learning using nonlinear dimensionality reduction (November 15, 2011). (Assigned International Business Machines Corporation) (Discussing discounted parameters.) *See also* EUGENE CHARNAK, INTRODUCTION TO DEEP LEARNING 113-115 (2018).

⁴⁹ EUGENE CHARNAK, INTRODUCTION TO DEEP LEARNING 114-115 (2018). *See also* U.S. Patent No. 8,060,454 to Das, et al., Method and apparatus for improved reward-based learning using nonlinear dimensionality reduction (November 15, 2011). (Assigned International Business Machines Corporation)

⁵⁰ MYKEL J. KOCHENDERFER, DECISION MAKING UNDER UNCERTAINTY 80 (2015).

⁵¹ Policy iteration is a method of finding the optimal policy by continuously evaluating and improving the policy. *See* U.S. Patent No. 9,661,019 to Liu, System and method for distributed denial of service identification and prevention (May 23, 2017). (Assigned to Oracle International Corporation)

⁵² MYKEL J. KOCHENDERFER, DECISION MAKING UNDER UNCERTAINTY 80 (2015). *See also* U.S. Patent No. 8,468,041 to Vengerov, Using reinforcement learning to facilitate dynamic resource allocation (June 18, 2013) (Assigned to Oracle America, Inc.)

⁵³ MYKEL J. KOCHENDERFER, DECISION MAKING UNDER UNCERTAINTY 81 (2015).

⁵⁴ U.S. Patent No. 9,661,019 to Liu, System and method for distributed denial of service identification and prevention (May 23, 2017). (Assigned to Oracle International Corporation)

B. Deep Decisions

New technologies often represent a convergence of many different streams of techniques, devices, and machines, each coming from its own separate historical avenue of development.⁵⁵

Deep reinforcement learning is a new type of machine learning resulting from the technical convergence of two more mature machine learning methods, deep learning and reinforcement learning.⁵⁶ Deep reinforcement learning is important because it is a scalable method for general intelligence – a machine capable of achieving any definable goal.

In fact, deep reinforcement learning is also the most powerful type of AI. Deep reinforcement learning systems have three capabilities that set them apart from all previous systems: (1) generalization; (2) learning; and (3) intelligence. As such, attention is rapidly turning to deep reinforcement learning as the hottest area in AI research. For example, according to MIT Professor, Max Tegmark, “deep reinforcement learning is a completely general technique.”⁵⁷

Deep learning is a sub-field of machine learning concerned with the acquisition of knowledge from large amounts of data.⁵⁸ Deep learning involves modeling the human brain with machines to process information.⁵⁹ Deep learning is a process by which neural networks learn from large amounts of data.⁶⁰ The internet is the driving force behind modern deep learning

⁵⁵ PAUL E. CERUZZI, *COMPUTING: A CONCISE HISTORY* (2012).

⁵⁶ Amazing work in early state pre-processing is being done by Olga Russakovsky and Serena Yeung. See Olga Russakovsky, et al., *Best of both worlds: human-machine collaboration for object annotation* (2015), <https://ieeexplore.ieee.org/document/7298824>. (Introducing a model that integrates multiple computer vision models with multiple sources of human input in a Markov Decision Process.) Serena Yung, et al., *Every Moment Counts: Dense Detailed Labeling of Actions in Complex Videos* (2017) <https://arxiv.org/abs/1507.05738>. (Modeling multiple dense labels benefits from temporal relations within and across classes.)

⁵⁷ MAX TEGMARK, *LIFE 3.0: BEING HUMAN IN THE AGE OF ARTIFICIAL INTELLIGENCE* 83 (2017).

⁵⁸ ETHEM ALPAYDIN, *MACHINE LEARNING* 3 (2016). See also MICHAEL BUCKLAND, *INFORMATION AND SOCIETY* 21-22 (2017). (Discussing definitions of information.)

⁵⁹ See Michael Simon, et. al., *Lola v. Skadden and the Automation of the Legal Profession*, 20 *YALE J.L. & TECH.* 234, 254 (2018). See also ETHEM ALPAYDIN, *MACHINE LEARNING* 88-90 (2016).

⁶⁰ Brian S. Haney, *The Perils & Promises of Artificial General Intelligence*, 45 *J. LEGIS.* 151, 157 (2018). (Data are a digital representation of information about the world.)

strategies because the internet has enabled humanity to organize and aggregate massive amounts of data.⁶¹

Indeed, the explosion in data collection since the inception of the internet continues to result in increasingly available data, as well as improved deep learning applications and models.⁶² Critically, every day humans create five exabytes of data,⁶³ as much data as civilization created from the dawn of time until 1999.⁶⁴ This is particularly important because the data – not human programmers – drive progress in deep learning applications.⁶⁵ Generally, deep learning systems are developed in three parts: data pre-processing, model design, and learning. Systems integrating reinforcement learning and deep learning technologies, represent an unprecedented breakthrough for AI, showing state-of-the-art performance in nearly all industrial domains – from healthcare to modern warfare.⁶⁶

Deep Q-Networks (DQNs) are deep neural networks embedded in the reinforcement learning architecture, representing these two systems' convergence.⁶⁷ Indeed, the DQN is a critically important deep reinforcement learning algorithm.⁶⁸ The DQN algorithm develops an

⁶¹ RICHARD SUSSKIND, TOMORROW'S LAWYERS 11 (2017).

⁶² PETER J. DENNING, MATTI TEDRE, COMPUTATIONAL THINKING 93 (2019).

⁶³ An exabyte is 10^{18} or one quintillion byte.

⁶⁴ RICHARD SUSSKIND, TOMORROW'S LAWYERS 11 (2017).

⁶⁵ RICHARD SUSSKIND, TOMORROW'S LAWYERS 11 (2017).

⁶⁶ Serena Yeung, et al., A computer vision system for deep learning-based detection of patient mobilization activities in the ICU, *npj Digit. Med.* 2, 11 (2019). <https://doi.org/10.1038/s41746-019-0087-z>. (Introducing an algorithm for detection of mobility activity occurrence.)

⁶⁷ Leslie Pack Kaelbling, *Learning in Embedded Systems* (1990), <https://apps.dtic.mil/dtic/tr/fulltext/u2/a323936.pdf>.

⁶⁸ U.S. Patent No. 10,032,281, Multi-scale deep reinforcement machine learning for N-dimensional segmentation in medical imaging (July 24, 2018), Ghesu, et al. (Assigned to Siemens Healthcare) *See also* Yuval Tassa, et. al., *DeepMind Control Suite*, 12 (January 3, 2018) (<https://arxiv.org/abs/1801.00690>). (The Deep Mind Control Suit is a set of tasks for benchmarking continuous RL algorithms developed by Google Deep Mind.) *See also* U.S. Patent No. 10,296,830, to Cai, et al. Dynamic topic guidance in the context of multi-round conversation (May 21, 2019). (Assigned to International Business Machines Corporation).

optimal policy⁶⁹ for an agent with a Q-learning algorithm.⁷⁰ Q-Learning is a model-free reinforcement learning technique, a trial-and-error algorithm.⁷¹ NYU Law Professor Jeanne Fromer explains, “Q-learning seems especially suited for learning the most successful actions in a particular state for a system.”⁷²

The DQN algorithm combines Q-learning⁷³ with a neural network to maximize an agent’s reward.⁷⁴ The DQN algorithm’s most important aspect is the Bellman Equation.⁷⁵ The Bellman Equation does two things; it defines the optimal policy and forces the agent to consider the

⁶⁹ The optimal policy is the best method of decision making for an agent with the goal of maximizing reward. *See* MYKEL J. KOCHENDERFER, *DECISION MAKING UNDER UNCERTAINTY* 81 (2015).

⁷⁰ Serena Yeung, et al., *Learning to Learn from Noisy Web Videos* (2017), http://openaccess.thecvf.com/content_cvpr_2017/html/Yeung_Learning_to_Learn_CVPR_2017_paper.html. (Using Q-learning to learn a data labeling policy on a small labeled training dataset, and then uses this to automatically label noisy web data for new visual concepts.)

⁷¹ Leslie Pack Kaelbling, et al., *Reinforcement Learning: A Survey*, *J. of Artificial Intelligence Research* 253 (1996), <http://www.cse.msu.edu/~cse841/papers/kaelbling.pdf>.

⁷² Jeanne C. Fromer, *Learning Optimal Discourse Strategies in a Spoken Dialogue System*, MIT Master’s Thesis, 40, 43 (1998), <https://dspace.mit.edu/bitstream/handle/1721.1/47703/42306186-MIT.pdf?sequence=2&isAllowed=y>. (“Q-learning seems especially suited for learning the most successful actions in a particular state for a system.” And, further explaining, “These algorithms can calculate optimal discourse policies for Markov decision problems (MDPs), accessible, stochastic environments with a known transition model.”)

⁷³ United States Patent No. 10,049,301 to Kluckner, et al., *Medical scanner teaches itself to optimize clinical protocols and image acquisition* (August 14, 2018). (Assigned to Siemens Healthcare.) *See also* MAXIM LAPAN, *DEEP REINFORCEMENT LEARNING HANDS-ON* 144 (2018).

⁷⁴ PAUL JOHN WERBOS, *THE ROOTS OF BACKPROPAGATION FROM ORDERED DERIVATIVES TO NEURAL NETWORKS AND POLITICAL FORECASTING* 306-307 (1994). *See also* Jeanne C. Fromer, *Learning Optimal Discourse Strategies in a Spoken Dialogue System*, MIT Masters Thesis, 43 (1998), <https://dspace.mit.edu/bitstream/handle/1721.1/47703/42306186-MIT.pdf?sequence=2&isAllowed=y>. (“Q-learning seems especially suited for learning the most successful actions in a particular state for a system.”)

⁷⁵ In short, the Bellman Equation expresses the relationship between the value of a state and the values of its successor states.

$$Q * (S_t, a_t) = E[r_t + Ymax_a Q(S_{t+1}, a)]$$

The algorithm continues perpetually until the Q-value function’s convergence with an approximate maximum.

$$Q'_j \rightarrow Q'_{j+1} \rightarrow Q'_{j+2} \rightarrow \dots \rightarrow Q *$$

See U.S. Patent No. 8,060,454 to Das, et al., *Method and apparatus for improved reward-based learning using nonlinear dimensionality reduction* (November 15, 2011). (Assigned International Business Machines Corporation) (Claim 14 and claim 23 both discuss applications of Bellman equations for optimality.)

reward in its present state as greater compared to rewards in future states.⁷⁶ However, the Bellman Equation is a slower algorithm in practice and can be computationally expensive. Thus, a neural network is used as an approximator for a state-action value function, allowing for more efficient programming and model development.⁷⁷

The DQN is an off-policy algorithm, meaning it uses data to optimize performance.⁷⁸ However, one problem with training DQN algorithms, and off-policy deep reinforcement learning algorithms more generally, is that they are computationally expensive. In other words, these algorithms require massive computing power.⁷⁹ As such, these system's intersection with quantum computers provides a better way to develop AI technology.⁸⁰ Still, reinforcement learning applications for real world problem solving are scaling the economy.

⁷⁶ Jeanne C. Fromer, Learning Optimal Discourse Strategies in a Spoken Dialogue System, MIT Masters Thesis, 43 (1998), <https://dspace.mit.edu/bitstream/handle/1721.1/47703/42306186-MIT.pdf?sequence=2&isAllowed=y>. (“Q-learning seems especially suited for learning the most successful actions in a particular state for a system.”)

⁷⁷ Q-value correlates with an optimal policy, telling the agent which actions to take in any given state. See Volodymyr Mnih, Koray Kavukcuoglu, Methods and Apparatus for Reinforcement Learning, U.S. Patent Application No. 14/097,862 at 5 (filed Dec. 5, 2013), <https://patents.google.com/patent/US20150100530A1/en>.

⁷⁸ Hado van Hasselt, Arthur Guez, and David Silver, *Deep Reinforcement Learning with Q-Learning*, Google DeepMind, 2098 (2018) <https://arxiv.org/abs/1509.06461>.

⁷⁹ Brian S. Haney, AI Patents: A Data Driven Approach, 19 CHI.-KENT J. INTELL. PROP. 407, 441 (2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3527154.

⁸⁰ Brian S. Haney, *Quantum Machine Learning: A Patent Review*, 12 Case W. Res. J.L. Tech. & Internet __ (2020) (Forthcoming), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3626534.

C. Applied Art

Reinforcement learning applications cut vertically across all sectors in the information economy. Reinforcement learning systems propel the state-of-the-art in transportation, defense, healthcare,⁸¹ information technology,⁸² and cybersecurity.⁸³

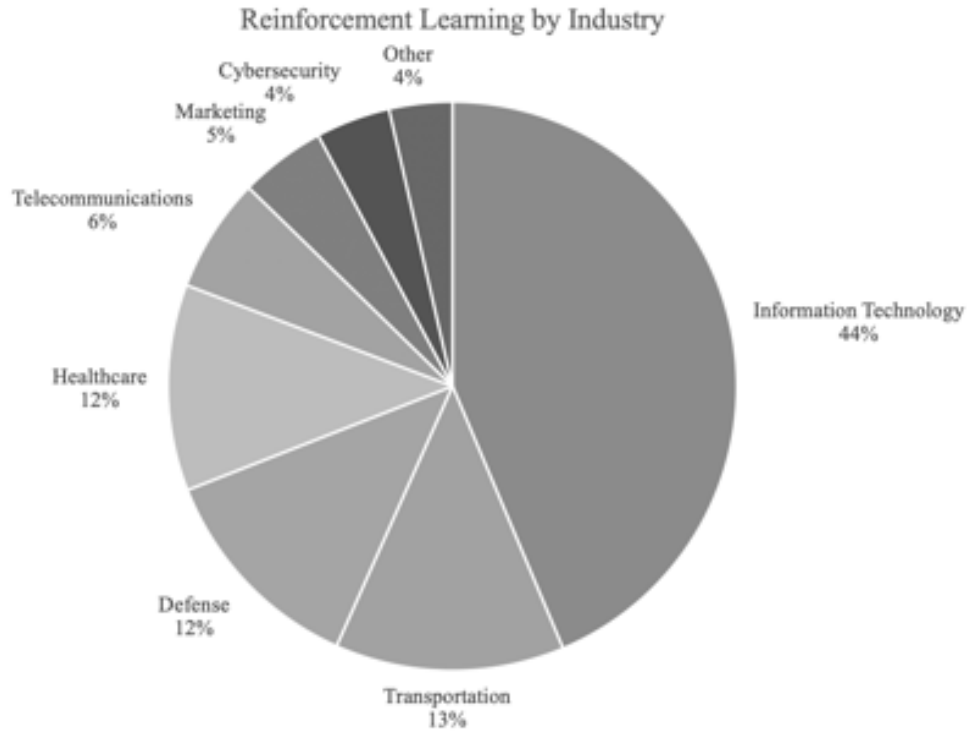


Figure 3⁸⁴

⁸¹ Robert A. Edgell, Roland Vogl, *A Network View of Human Ingestion: Instrumental Artificial Intelligence*, Artificial Intelligence and Smarter Living Workshop, AAAI Spring Symposium, 25 (2011), <https://www.aaai.org/ocs/index.php/WS/AAAIW11/paper/viewFile/3886/4214>. (“Although our initial review of the literature suggests that AI could instrumentally mediate health outcomes, many questions remain, especially in the social realm. In addition to the direct informational and cognitive challenges outlined above, other stakeholder difficulties may emerge.”) See also U.S. Patent No. 10,032,281, Multi-scale deep reinforcement machine learning for N-dimensional segmentation in medical imaging (July 24, 2018). See also EP3399501A1, Multi-Scale Deep Reinforcement Machine Learning for N-Dimensional Segmentation in Medical Imaging (November 7, 2018).

⁸² U.S. Patent 10,546,066, End-to-end learning of dialogue agents for information access (January 28, 2020).

⁸³ US 2019311119 A1, Deep Reinforcement Learning Technologies for Detecting Malware to Mihail, et al., (October 10, 2019). See also WO2019226402A1, Training Technologies for Deep Reinforcement Learning Technologies for Detecting Malware to Mihail, et al. (November 28, 2019). See also EP 3,528,463 A1, An artificial intelligence cybersecurity analyst (August 21, 2019).

⁸⁴ Transatlantic Reinforcement Learning Patent Dataset.

Figure 3 depicts the industrial breakdown for reinforcement learning patents. By far, the most common application for reinforcement learning is information technology. This is unsurprising given reinforcement learning methods are at the bleeding edge in machine learning optimization.⁸⁵ Top firms have patents in this area, including Amazon and Oracle.⁸⁶

IBM has a dominant stakeholder in the machine learning patent market generally.⁸⁷ Their presence in the transatlantic reinforcement learning market is no exception.




Figure 4⁸⁸

⁸⁵ Maria Schuld, et al., Prediction by linear regression on a quantum computer, 1 (2016), <https://arxiv.org/abs/1601.07823v2>. (“The central problem of machine learning is pattern recognition, in which a machine is supposed to infer from a set of training data how to map new inputs of the same type to corresponding outputs.”) See also Katerina Fragkiadaki, et al., Figure-Ground Image Segmentation helps Weakly-Supervised Learning of Objects (2010) https://link.springer.com/chapter/10.1007/978-3-642-15567-3_41. (Optimizing a conditional likelihood of the image collection given the image bottom-up saliency information.)

⁸⁶ U.S. Patent No. 8,429,096 to Soundararajan, et al. Resource isolation through reinforcement learning (April 23, 2013). (Assigned to Amazon) See also U.S. Patent No. 8,468,041 to Vengerov, Using reinforcement learning to facilitate dynamic resource allocation (June 18, 2013). (Assigned to Oracle)

⁸⁷ Brian S. Haney, AI Patents: A Data Driven Approach, 19 CHI.-KENT J. INTELL. PROP. 407, 479 (2020).

⁸⁸ Transatlantic Reinforcement Learning Patent Dataset.

Figure 4 graphs the number of reinforcement learning patents by firm. Top information technology, telecommunications, and transportation firms are noticeably present in the market.

II. Patent Data

One of the most difficult problems with developing scalable AI is tracking innovation.⁸⁹ As such, this research will provide empirical data relating to a keystone architecture in the machine learning paradigm. Data was collected from both the USPTO and the EPO for patents with claim terms for reinforcement learning thought December 31, 2019. The USPTO lists 234 U.S. Patents for reinforcement learning. And, the EPO currently lists 90 patents in its database for patents with claims for reinforcement learning.

Transatlantic Reinforcement Learning Patents

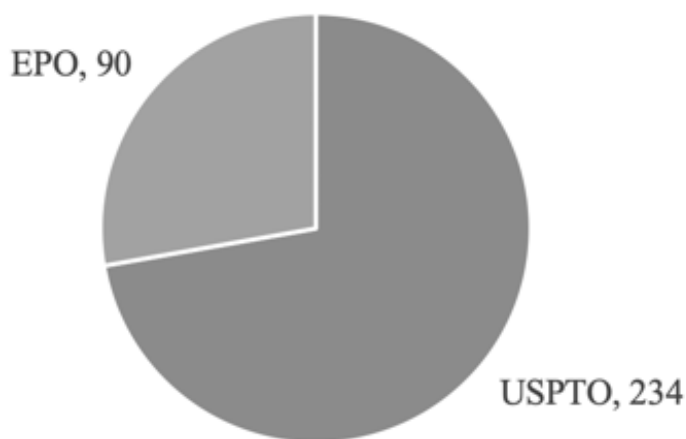


Figure 5⁹⁰

⁸⁹ One reason is because programs developing AI are segmented across countless Government agencies and private firms. See generally Veronica Root Martinez, *Complex Compliance Investigations*, 120 COLUMBIA L. R. 249, 274 (2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3350463. (“Importantly, examples of silos abound both within private organizations and the government.”)

⁹⁰ Brian Haney, *Transatlantic Reinforcement Learning Patents* (2020). (Data on file with author.)

This Part aggregates and analyzes reinforcement learning patent data from both the USPTO and EPO to provide a transatlantic perspective on reinforcement learning innovation and intellectual property protections.

The analysis was performed with the Python Pandas library,⁹¹ to calculate value properties for the U.S. Patents in the Transatlantic Reinforcement Learning Patent Dataset. Data for each patent element was gathered in a separate file and subsequently processed as input to a computer program, which produces statistical results. The analysis focused on three patent metrics associated with patent value, claims, prior art citations, and inventors.

A. USPTO

Figure 6 displays analysis performed on the U.S. Patents in the Transatlantic Reinforcement Learning Patent Dataset.

	Claims	Prior Art	Inventors
Count	234.00	234.00	234.00
Mean	21.27	44.44	2.86
Mode	20.00	10.00	1.00
Minimum	1.00	2.00	1.00
Quartile 1	17.00	11.00	1.00
Quartile 2	20.00	21.00	2.00
Quartile 3	23.75	38.00	4.00
Maximum	94.00	753.00	13.00

*Figure 6*⁹²

In total, there were 234 U.S. patents included in the dataset. U.S. Patent No. 6,572,542, System and method for monitoring and controlling the glycemic state of a patient, which was awarded to Medtronic, Inc. had the most claims for any single patent with 94.⁹³ U.S. Patent No. 9,948,788,

⁹¹ Pandas, Project Description (2020), <https://pypi.org/project/pandas/>.

⁹² Brian Haney, Reinforcement Learning Patents (2020). On file with author.

⁹³ U.S. Patent No. 6,572,542, System and method for monitoring and controlling the glycemic state of a patient (June 3, 2003).

Method and system for preventing illicit use of a telephony platform, which was awarded to Twilio cited the most prior art with 753 citations.⁹⁴

B. EPO

Figure 7 displays analysis performed on the European Patents in the Transatlantic Reinforcement Learning Patent Dataset.

	Claims	Prior Art	Inventors
Count	90.00	90.00	90.00
Mean	18.24	5.19	3.02
Mode	15.00	4.00	1.00
Minimum	9.00	0.00	1.00
Quartile 1	14.00	2.00	2.00
Quartile 2	15.00	4.00	3.00
Quartile 3	18.75	6.00	4.00
Maximum	98.00	89.00	12.00

Figure 7⁹⁵

In total, there were 90 EPO patents included in the dataset. European Patent No. 1,528,464 Proactive user interface including evolving agent, has the most claims at 98 and is assigned to Samsung.⁹⁶ European Patent No. 2,325,294 Visual-Servoing Optical Microscopy, cited the most prior art at 89 and is assigned to the University of California.⁹⁷

C. Transatlantic

Figure 8 displays analysis performed on the Transatlantic Reinforcement Learning Patent Dataset.

⁹⁴ U.S. Patent No. 9,948,788, Method and system for preventing illicit use of a telephony platform (April 17, 2018).

⁹⁵ Brian Haney, Reinforcement Learning Patents (2020). On file with author.

⁹⁶ EP1528464A2, Proactive user interface including evolving agent (May 4, 2005).

⁹⁷ EP No. 2,325,294 Visual-Servoing Optical Microscopy (May 25, 2011).

	Claims	Prior Art	Inventors
Count	324.00	324.00	324.00
Mean	20.44	33.54	2.91
Mode	20.00	4.00	1.00
Minimum	1.00	0.00	1.00
Quartile 1	15.00	5.00	1.00
Quartile 2	19.00	13.00	2.00
Quartile 3	22.00	30.00	4.00
Maximum	98.00	753.00	13.00

Figure 8⁹⁸

In total, there were 324 patents included in the dataset. U.S. Patents tended to include more claims with a 21.27 mean compared to European Patents, which had a mean of 18.24. The most noticeable difference was in the number of prior art citations. The mean for U.S. Patents was 44.44 prior art citations per patent compared to 5.19 for European Patents. The average number of inventors for both European and U.S. Patens was about three, which is reflected in the 2.91 average for the Transatlantic Dataset.

III. Comparative Claims Critique

Claims are the patent’s most important feature because claims are the only part that can be infringed. This part will compare and contrast claims from the USPTO and EPO. First, analysis is provided for machine translation techniques among transatlantic languages. The three main languages are French, German, and English. In fact, the EPO translates patents across these three languages. Second, the syntactic differences between USPTO and EPO patents will be examined with reference to comparative analysis for reinforcement learning patent claims.⁹⁹

⁹⁸ Brian Haney, Reinforcement Learning Patents (2020). On file with author.

⁹⁹ Anna Maria Baumgartner, *Antitrust Issues in Technology Transfer: A Comparative Legal Analysis of Patent Licenses in the EU and the U.S.*, TTLF Working Papers 145 (2013), <https://law.stanford.edu/publications/antitrust-issues-in-technology-transfer-a-comparative-legal-analysis-of-patent-licenses-in-the-eu-and-the-u-s/>. (“Case law in both jurisdictions has established that IP rights are neither free from antitrust scrutiny nor particularly suspect.”)

Third, transatlantic patent scope issues will be addressed. Scope is important for complex concepts like reinforcement learning because subject matter expertise rests on a high curve. As a result, patent thickets¹⁰⁰ are spawning in the art.¹⁰¹ The critique will contribute a transatlantic perspective, discussing Reinforcement Learning Patents granted by both the USPTO and EPO.

A. Technically Translating

The Ancients believed there was a time at which the whole world had a common language.¹⁰² Perhaps, they were right and similar to the way in which entropy guides our Universe from order to disorder,¹⁰³ time drives the expansion of language from cogeny to chaos.¹⁰⁴ Regardless, “The epigenesis of our highly articulated human language is a fascinating story.”¹⁰⁵

There is an inherent difficulty in defining quality metrics for translation¹⁰⁶ because at the heart of natural language is vagueness and ambiguity.¹⁰⁷ Even more difficult is the necessity for

¹⁰⁰ Mark Lemley and Carl Shapiro, Probabilistic Patents (2004), <http://ssrn.com/abstract=567883>. (The result is a “patent thicket” in which hundreds of patents overlap with each other, creating a barrier that is hard for a competitor to penetrate.)

¹⁰¹ Anna Maria Baumgartner, *Antitrust Issues in Technology Transfer: A Comparative Legal Analysis of Patent Licenses in the EU and the U.S.*, TTLF Working Papers 145 (2013), <https://law.stanford.edu/publications/antitrust-issues-in-technology-transfer-a-comparative-legal-analysis-of-patent-licenses-in-the-eu-and-the-u-s/>. (“Moreover, neither in the EU nor in the U.S. are intellectual property rights presumed to confer market power.”)

¹⁰² According to the bible, humans were incredibly powerful speaking the same language, deciding to build a bridge to the heavens called the Tower of Babel. But then God said, “If as one people speaking the same language, they have begun to do this, then nothing they plan to do will be impossible for them. Come, let us go down and confuse their language so they will not understand each other.” So, God scattered the people’s language across the Earth, stopping the Tower of Babel’s construction. Genesis 11:2-11, *See also* Otto Jespersen, *Language: Its Nature Development and Origin* (1922). (“Aristotle laid the foundation of the division of words into “parts of speech” and introduced the notion of case.”)

¹⁰³ BRIAN GREENE, *FABRIC OF THE COSMOS* 151 (2005).

¹⁰⁴ Frederic H. Behr, Jr., et al., Estimating and Comparing Entropy Across Written Natural Languages using PPM Compression, Harvard Computer Science Group (2017) <https://www.irif.fr/~dxiao/docs/entropy.pdf>. *See also* Daniel Martin Katz, et al., Legal N-Grams? A Simple Approach to Track the ‘Evolution’ of Legal Language (2011) https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1971953.

¹⁰⁵ ZOLTAN TOREY, *THE CONSCIOUS MIND* 47 (2014). (Epigenesis describes a theory of development through gradual differentiation.)

¹⁰⁶ Thierry Poibeau, *Machine Translation* 8 (2017).

¹⁰⁷ NOAM CHOMSKY, *SYNTACTIC STRUCTURES* 17 (1957).

understanding context for semantic derivation.¹⁰⁸ Semantic derivation refers to the process by which syntax is distilled to meaning. Moreover, translators must understand both the textual language and the target language.¹⁰⁹ A translator must know both languages and have the skill to reformulate a source language in a target language.¹¹⁰ Further, word senses are not mutually exclusive and “word usages often fall between dictionary definitions.”¹¹¹

Despite these difficulties, research and development efforts for machine translation rage on. The idea for machine translation is not to replace human translators who are the only ones able to translate novels or poetry.¹¹² The overall quality achievable by machine translation has also been a matter of much debate.¹¹³ The ultimate goal is to obtain a quality translation equivalent to that of a human being.¹¹⁴

Bayesian statistics¹¹⁵ reverses the problem and aims to determine, given various sequences from the target language, which has the most chance of being a translation of the source language.¹¹⁶ The statistical translation process can be decomposed into three steps: (1)

¹⁰⁸ Thierry Poibeau, Machine Translation 18 (2017).

¹⁰⁹ Thierry Poibeau, Machine Translation 10-11 (2017). (“What is expected of a translation can vary radically depending on the clients, the era, the nature of the text, its usage, or even context.”)

¹¹⁰ Thierry Poibeau, Machine Translation 11-12 (2017). (“Word-for-word translation is not good practice, since the result is often hard to understand and not idiomatic in the target language.”)

¹¹¹ Thierry Poibeau, Machine Translation 18 (2017).

¹¹² Thierry Poibeau, Machine Translation 13 (2017).

¹¹³ Thierry Poibeau, Machine Translation 13 (2017).

¹¹⁴ Thierry Poibeau, Machine Translation 13 (2017).

¹¹⁵ Fang Liu, Assessment of Bayesian Expected Power via Bayesian Bootstrap 14 (2017),

<https://arxiv.org/abs/1705.04366>. (“The bootstrap-based procedures will appeal to non-Bayesian practitioners given their analytical and computational simplicity and easiness in implementation.”) *See also* Ava P. Soleimany, et al., Image Segmentation of Liver Stage Malaria Infection with Spatial Uncertainty Sampling, 1 (2019),

<https://arxiv.org/abs/1912.00262>. (“To address this need for automation in the anti-malarial screening pipeline, we present a convolutional neural network-based (CNN) architecture for automated segmentation of parasites in liver stage malaria infection and develop a Bayesian deep learning approach for estimating uncertainty in image segmentations.”) *See also* Maria Schuld, et al., An introduction to quantum machine learning 2 (2014), <https://arxiv.org/pdf/1409.3097.pdf>. (Stochastic models such as Bayesian decision theory or hidden Markov models find an elegant translation into the language of open quantum systems.)

¹¹⁶ Thierry Poibeau, Machine Translation 127-128 (2017). This can be formalized:

$$Pr(T|S) = \frac{Pr(T)Pr(T|S)}{Pr(S)}$$

Determine the length of the target sentence depending on the length of the source sentence; (2) Identify the best possible alignment between the source sentence and the target sentence; and then (3) Find correspondences at word level.¹¹⁷ This strategy clearly gives a very simplified picture of the translation process.¹¹⁸

For example, machine learning scholar Ethem Alpaydin argues the driving force of computing technology is the realization that every piece of information can be represented as numbers.¹¹⁹ It follows logically that all information can be processed with computers. However, the divide between syntax and semantics – manifested in a computer’s inability to understand human language – remains one of the most challenging problems in artificial intelligence. Deceptively challenging tasks fall under the rubric of common-sense reasoning, a term encompassing different types of cognitive skills and real-world knowledge.¹²⁰

One important market¹²¹ for machine translation is access to international patents.¹²² Patents can be written in a wide variety of languages.¹²³ There is a specific need to break the language barriers in patents because a patent in a specific language can be the source of major problems with high financial impact.¹²⁴ Another issue is that patents are written in a very specific

The equation may be simplified:

$$T' = \operatorname{argmax}_T [Pr(T) * Pr(S|T)]$$

Arguably, this is the fundamental equation for machine translation.

¹¹⁷ Thierry Poibeau, Machine Translation 129 (2017).

¹¹⁸ Thierry Poibeau, Machine Translation 129 (2017).

¹¹⁹ ETHEM ALPAYDIN, MACHINE LEARNING 2 (2016).

¹²⁰ Brian S. Haney, Patents for NLP Software: An Empirical Review, 22 N.C. J.L. & Tech. __ (2020) (Forthcoming), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3594515.

¹²¹ The oldest and most well-known machine translation company is Systran. The Air Force was one of the company’s first customers and was interested in Russian-English translation. Then, it won a contract with the European Commission, which subsequently bid out the contract to other companies. Systran filed a lawsuit for copyright infringement and unauthorized data disclosure. Systran won the case in 2010. *See* Thierry Poibeau, Machine Translation 222-235 (2017).

¹²² Thierry Poibeau, Machine Translation 222 (2017).

¹²³ Thierry Poibeau, Machine Translation 222 (2017).

¹²⁴ Thierry Poibeau, Machine Translation 224 (2017).

jargon.¹²⁵ What is already difficult in one language only becomes more crucial and difficult in the context of multilingual systems.¹²⁶ The field of multilingual patents has drawn major investment due to the potential for high profits and risk for high loss.¹²⁷ For example, Google is working with European Patent Office to develop machine translation systems based on a statistical model originally developed by IBM.¹²⁸ And, the patents in this market are overflowing with thickets and overlap.

Additionally, the World Intellectual Property Organization is developing a translation system based on neural networks to translate Chinese, Japanese, and Korean documents to English.¹²⁹ As such, modern linguistic evolutions transcend intellectual property and innovation. According to Torey, the mind is made of concepts, which we express through words.¹³⁰ As such, our ability to formulate and describe new ideas is inherently limited by the words in our vocabulary. However, language need not be cage for the imagination.¹³¹ By cutting across languages in the innovation process, language may be a liberating passport to an open world.

B. Simplifying Syntax

Syntactic claim structures are a cogent way to paint an invention with a textual corpus.¹³² Often, claim construction is complicated by intentionally confusing analysis.¹³³ Motivated by

¹²⁵ Thierry Poibeau, Machine Translation 224-225 (2017).

¹²⁶ Thierry Poibeau, Machine Translation 225 (2017).

¹²⁷ Thierry Poibeau, Machine Translation 225 (2017).

¹²⁸ Thierry Poibeau, Machine Translation 228 (2017).

¹²⁹ Thierry Poibeau, Machine Translation 225 (2017).

¹³⁰ ZOLTAN TOREY, THE CONSCIOUS MIND 40 (2014).

¹³¹ ZOLTAN TOREY, THE CONSCIOUS MIND 47 (2014). (“[I]anguage need not be a cognitive trap but can become a liberating passport to ever-deepening insights into the world and conscious mind itself.”)

¹³² Burk, Dan L. and Lemley, Mark A., Quantum Patent Mechanics (2005). Lewis & Clark Law Review, Vol. 9, p. 29, 30 (2005) Stanford Public Law Working Paper No. 102, <https://ssrn.com/abstract=628224>. (“Determining that meaning necessarily requires the judge to break the text of a claim into discrete “elements” or units of text corresponding to the elements or units that comprise the claimed invention—essentially, organizing the language of the claims into “chunks” or “quanta” of text.”)

¹³³ Julie E. Cohen and Mark A. Lemley, Patent Scope and Innovation in the Software Industry, 89 Cal. L. Rev. 1, 37 (2001). (“In theory, this process of claim construction determines the scope of the patent. In practice, however, the doctrine of equivalents is the primary tool available to courts and litigants for fine-tuning patent scope.”)

monetary gain, firms conflate claim meaning and scope to design around and reverse engineer competitors technology.¹³⁴ But don't be fooled by complexity, simplicity is truth.¹³⁵

Every claim has three parts: (1) a preamble, (2) a transitional phrase, and (3) a body. The preamble describes the invention type, the transitional phrase sets the scope,¹³⁶ and the body defines the invention. Consider the following claim for a reinforcement learning application in cybersecurity:

“1. A method performed on at least one computing device that includes at least one processor and memory, the method comprising: training, by the at least one computing device, a deep reinforcement learning (“DRL”) model, where the training is based on a set of training files, where each training file of the set is associated with a label that indicates whether the each training file is considered malicious or benign, and where the training comprises: processing, by the DRL model from each file of the set of training files, a plurality of event states, where each event state comprises an event histogram, and where the processing further comprises considering the label of the each file; executing, by the at least one computing device, at least a portion of a file; and halting, by the at least one computing device in response to a decision by the trained DRL model, the execution of the at least the portion of the file.”¹³⁷

Here, the preamble is “A method performed on at least one computing device that includes at least one processor and memory.” The transitional phrase is “the method comprising” and the body is the remainder.

¹³⁴ Julie E. Cohen and Mark A. Lemley, Patent Scope and Innovation in the Software Industry, 89 Cal. L. Rev. 1, 20 (2001). (“Whether reverse engineering infringes a patent will further depend on the way the claim is written.”) See also Mark A. Lemley & Mark P. McKenna, Scope, 57 WM. & MARY L. REV 2197, 2240 (2015), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2660951. (“The patent claims themselves are an effort to define the scope of the legal right *ex ante*.”)

¹³⁵ *simplex sigillum veri* – simplicity is the sign of truth.

¹³⁶ *Comprising* means the invention includes but is not limited to the elements identified in the claim. *Consisting essentially of* limits the scope of a claim to the specified materials or steps and those that do not materially affect the basic and novel characteristics of the claimed invention. See *Mars Inc. v. H.J. Heinz Co.*, 377 F.3d 1369, 1376 (Fed. Cir. 2004) (“like the term ‘comprising,’ the terms ‘containing’ and ‘mixture’ are open-ended.”). See also *Genentech, Inc. v. Chiron Corp.*, 112 F.3d 495, 501, (Fed. Cir. 1997). (“Comprising” is a term of art used in claim language which means that the named elements are essential, but other elements may be added and still form a construct within the scope of the claim.)

¹³⁷ WO2019226402A1, Training Technologies for Deep Reinforcement Learning Technologies for Detecting Malware to Mihail, et al. (November 28, 2019).

Lord Polonius offers the ultimate syntactic truth, “brevity is the soul of wit.”¹³⁸ Not much has changed in the four centuries since Shakespeare scribed Hamlet. Stalled by technical stagnation, the past few centuries have yielded no real progress for human civilization. But perhaps new AI applications using reinforcement learning offer hope. For example, consider the following claim for a new radar system:

“1. An adaptive radar system (1) comprising: a transmitter unit (2) adapted to transmit a radar signal; a receiver unit (3) adapted to receive a reflected radar signal; and a reinforcement learning engine (4) adapted to compute a control policy (CP) used to control signal parameters and waveforms of the radar signal transmitted by the transmitter unit (2).”¹³⁹

Here, the preamble is “An adaptive radar system” and the transitional phrase is “comprising.” The claim body is relatively concise, reflecting more value and coverage.

In more competitive domains, like healthcare, claim syntax is more convoluted. Consider claim 9 from, European Patent No. 3,618,080, Control method and reinforcement learning for medical system:

“9. A medical system, comprising: an interaction system, configured for receiving an initial symptom; a decision agent interacting with the interaction system; and a neural network model, utilized by the decision agent to select at least one symptom inquiry action according to the initial symptom; wherein the interaction system is configured to receive at least one symptom answer in response to the at least one symptom inquiry action, wherein the neural network model is utilized by the decision agent to select at least one medical test action from candidate test actions according to the initial symptom and the at least one symptom answer, wherein the interaction system is configured to receive at least one test result of the at least one medical test action, and wherein the neural network model is utilized by the decision agent to select a result prediction action from candidate prediction actions according to the initial symptom, the at least one symptom answer and the at least one test result.”¹⁴⁰

¹³⁸ William Shakespeare, Hamlet, Scene 2 Act 2 (1600).

¹³⁹ EP 3,339,880 A1, Adaptive Radar System (June 27, 2018).

¹⁴⁰ EP 3,618,080 A1, Control method and reinforcement learning for medical system (April 3, 2020).

Now compare and contrast claim 9 from the 080' patent with claim 1, from United States Patent No. 10,032,281, Multi-scale deep reinforcement machine learning for N-dimensional segmentation in medical imaging:

“1. A method for three-dimensional segmentation based on machine learning in a medical imaging system, the method comprising: scanning, by a magnetic resonance, computed tomography, x-ray, or ultrasound imaging system, a patient, the scanning providing a medical dataset representing a three-dimensional region of a patient, the medical dataset comprising magnetic resonance, computed tomography, x-ray, or ultrasound data; applying, by a machine, the medical dataset to a multi-scale deep reinforcement machine-learned model, the multi-scale deep reinforcement machine-learned model trained with multi-scale deep reinforcement learning to segment boundaries of a three-dimensional object from the medical dataset, the multi-scale deep reinforcement machine-learned model including a machine-learned policy of a sequence of actions for shape evolution over iterative refinements of the boundaries of the three-dimensional object, the machine-learned policy trained to select each of the actions from possibilities given a state of fitting at each of the iterative refinements of shape representation parameters, the sequence resulting from the machine-learned policy training based on rewards, and the boundaries and refinements described using the shape representation parameters; rendering, by a renderer, an image of the three-dimensional object based on the boundaries determined with the policy; and displaying the image on a display device.”¹⁴¹

The first claim uses reinforcement learning in an interactive system for symptom analysis. By contrast, the second claim uses reinforcement learning for processing medical imaging data. But these claims have much more in common than differentiating features. For example, both claims incorporate neural networks, integrating reinforcement learning with deep learning technology.

Perhaps more significantly, both claims are shoddy syntactic work product and inventions, legal protections, and art. Critically, confusing claims do not happen by accident. Instead they result from an institutional process predicated on predatory deception. Moreover, claims with too many words are unduly drafted because they are too easy to design and patent around.

For example, if a firm wanted to patent around European Patent No. 3,618,080 and United States Patent No. 10,032,281, the firm may claim the following:

¹⁴¹ U.S. Patent No. 10,032,281, Multi-scale deep reinforcement machine learning for N-dimensional segmentation in medical imaging (July 24, 2018).

A method for medical diagnostics, the method comprising an information network collecting patient data relating to clinical symptoms through patient interview, technical scans, x-ray, ultrasound, or magnetic resonance, processing the data with a neural network, predicting patient diagnosis, and a reinforcement learning agent recommending treatment through a virtual display.

But perhaps design around is not so easy. For example, one may argue the preceding example is obvious.¹⁴² Maybe it is obvious deep reinforcement learning technology can optimize any function in accordance with a defined reward. But subjectively self-serving, obviousness depends on the artist's skill.

So, the innovative firm may design new syntax for reinforcement learning technologies, without even mentioning the term. For example:

A method for medical diagnostics, the method comprising a data machine mining clinical patient data through patient interview and digital harvesting, mixing the data with matrix simulations, forming a diagnosis, and recommending an optimal treatment or cure.

Most lawyers like to add words to claims because writing more words requires more time. And more time means more money. The patent office has incentive to push applicants to add more words because it presumably limits monopoly, the product for sale. This syntactical incentive structure, and the institution within, stops innovation. Like great inventions,¹⁴³ great claims do more with less.

¹⁴² Jeanne C. Fromer, *The Layers of Obviousness in Patent Law*, 22 HARV. J. OF L. & TECH., 75 (2008). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1119723. (“The nonobviousness doctrine seeks to ensure that patents are granted only for technologically significant advances to foster the patent system’s goal of stimulating useful innovation.”)

¹⁴³ See U.S. Patent No. 381,970 to Tesla, *System of Electrical Distribution* (May 1, 1888). See also U.S. Patent No. 265,786 to Edison, *Apparatus for The Transmission of Electrical Power* (October 10, 1882). See also U.S. Patent No. 219,268 to Edison, *Electric-Light* (September 16, 1879). See also U.S. Patent No. 401,520, *Method of Operating Electromagnetic Motors* (April 16, 1889).

C. Scaling Scope

A claim's scope defines the patent's protectable rights.¹⁴⁴ Most lawyers think scope is the keystone to patent valuation – more scope means more legal protection and freedom to act. To an extent this is correct because arguing with a patent examiner over scope is a time-consuming process, which means more capital value for lawyers. But as an asset to the firm or inventor, more patent scope probably means more cost and thus more liability.¹⁴⁵ Consider scope doesn't matter at all for reinforcement learning patents because no judge, jury, or examiner has the requisite skill in the art necessary to make an informed decision about the technology.¹⁴⁶

Or maybe the knowledge gap means scope is even more important for reinforcement learning because subject matter expertise rests on a high curve. Regardless, patent thickets¹⁴⁷ are spawning in this budding domain. Yet, there may be little legitimacy to patents in the \$49 billion market for reinforcement learning technologies.¹⁴⁸ Even within the United States Patent Office and European Patent Office, there are patented claims for legally contrasting subject matter.

In 2013 the United States Patent Office, awarded two patents for resource management using reinforcement learning. The first, U.S. Patent No. 8,429,096, *Resource isolation through reinforcement learning*, was awarded to Amazon in April.

¹⁴⁴ Mark A. Lemley & Mark P. McKenna, *Scope*, 57 WM. & MARY L. REV. 2197, 2209 (2015). (accessed at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2660951) (IP regimes require, not just similarity between the defendant's and plaintiff's works, but similarity with respect to the protectable elements.)

¹⁴⁵ Shawn P. Miller, *What's The Connection Between Repeat Litigation And Patent Quality? A (Partial) Defense of The Most Litigated Patents*, 16 Stan. Tech. L. Rev. 313, 317 (2013). (“Legal fees for one case can range from \$500,000 through summary judgment to over \$4 million through trial.”)

¹⁴⁶ Of course, that won't stop the decision from happening. See Julie E. Cohen and Mark A. Lemley, *Patent Scope and Innovation in the Software Industry*, 89 Cal. L. Rev. 1, 4 (2001). (“Thus, determining the scope of software patents will take us a long way towards determining what to do in practice with Internet business method patents as well.”)

¹⁴⁷ Mark A Lemley, *Software Patents and The Return of Functional Claiming* 2013 WIS. L. REV. 905, 906 (2013), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2117302. (Arguing software patents create “thickets” of overlapping inventions.)

¹⁴⁸ Calculated as a share in the AI market. See United States Government Accountability Office, *Technology Assessment, Artificial Intelligence* (2018). (Stating the global impact for AI is \$10.7 Trillion.) See also McKinsey Global Institute, *Notes from the AI Frontier*, (2019). (Stating the global impact for AI is \$10.7 Trillion.)

“1. A method, comprising: performing, by a computer: storing data on behalf of a plurality of subscribers in a shared storage system; receiving input specifying one or more service level parameter values, each parameter value specifying a quantity of a respective resource that can be allocated for use in executing one or more queries in the shared storage system; receiving a plurality of queries directed to the data; and executing the plurality of queries, wherein said executing comprises: monitoring resource utilization, wherein the monitored resource utilization is resource utilization by the plurality of queries of resources for which service level parameter values have been specified, wherein said monitoring is performed during execution of the plurality of queries; determining whether the resource utilization by the plurality of queries is consistent with the one or more specified service level parameter values; and in response to determining that the resource utilization by one or more of the plurality of queries is not consistent with the one or more specified service level parameter values, attempting to correct the determined inconsistency, wherein said attempting comprises applying a reinforcement learning technique to automatically update an execution parameter of at least one of the plurality of queries, wherein an execution parameter is a modifiable parameter affecting execution of the plurality of queries, and wherein the update is performed during execution of the plurality of queries.”¹⁴⁹

The first problem with this claim is the preamble does not include any description for the methods use. If an idea is not useful, it is not an invention. Regardless, the transitional phrase “comprising” indicates a broader reading, but the word count at 247 narrows the claim’s scope. Moreover, the syntax choice which includes, “plurality of queries”, “monitoring resource utilization” and “the determined inconsistency” probably renders the claim altogether indefinite.

Two months later the second patent, U.S. Patent No. 8,468,041 to Vengerov, *Using reinforcement learning to facilitate dynamic resource allocation*, was awarded to Oracle in June. Oracle’s patent claim is slightly better because it is more concise, clear, and cogent in form and substance.

“1. A method for allocating resources to projects in a computer system, comprising: in a resource allocation mechanism in the computer system, for each project, at least one processor of the computer system performing operations for: determining a current demand by the project for a resource; determining a current allocation of the resource to the project; using a computational model for the project to compute an expected long-term utility of the project for the resource; and using a reinforcement learning technique to update parameters of the computational model for the project based on performance

¹⁴⁹ U.S. Patent No. 8,429,096 to Soundararajan, et al. Resource isolation through reinforcement learning (April 23, 2013). (Assigned to Amazon)

feedback parameters; and trading the resource between the projects in the computer system to optimize a weighted summation of the computed expected long-term utilities of all the projects, wherein the weighted summation of the computed expected long-term utilities is a measure of a single common benefit of all the projects.”¹⁵⁰

The preamble phrase “for allocating resources to projects in a computer system” illustrates useful purpose. As for scope, the transitional phrase is comprising which denotes a broader reading. However, similar to the Amazon patent, the syntactic choice for the claim body is atrociously indefinite. Critically, the claim is written in passive voice “processor of the computer”, “parameters of the computational model”, and “utilities of all the projects.” The passive voice reflects a passive claim.

The European Patent Office faces similar difficulties. For example, consider two patents for target-based processing using reinforcement learning. First, European Patent 3,553,711, Information Processing Device and Method, and Program, is assigned to Sony.

“1. An information processing apparatus comprising: a reward calculating part configured to calculate a reward based on an input objective state of a control target and a state of the control target based on a sensing result of the control target; a reinforcement learning part configured to perform reinforcement learning using the reward calculated by the reward calculating part and the state of the control target to select a better action for bringing the control target closer to the objective state; and an action execution part configured to execute the action selected by the reinforcement learning part for the control target.”¹⁵¹

The preamble is strong, clearly defining the invention type an “information processing apparatus.” The transitional term “comprising” and relatively limited word count 115, support a broader construction. The body describes the invention well, an apparatus using reinforcement learning for target control.

¹⁵⁰ U.S. Patent No. 8,468,041 to Vengerov, Using reinforcement learning to facilitate dynamic resource allocation (June 18, 2013). (Assigned to Oracle)

¹⁵¹ European Patent 3553711A1, Information Processing Device and Method, and Program (2019). (Assigned to Sony)

Second consider, European Patent 3,129,839, Controlling a target system, which is assigned to Siemens. The patent uses a list to describe the invention as a process.

“1) A method for controlling a target system by a processor on the basis of a pool of control policies, the method comprising : a) receiving the pool of control policies comprising a plurality of control policies, b) receiving weights for weighting each of the plurality of control policies, c) weighting the plurality of control policies by the weights to provide a weighted aggregated control policy, d) controlling the target system using the weighted aggregated control policy, e) receiving performance data relating to a performance of the controlled target system, f) adjusting the weights by the processor on the basis of the received performance data to improve the performance of the controlled target system, and g) reweighting the plurality of control policies by the adjusted weights to adjust the weighted aggregated control policy.”¹⁵²

This is arguably helpful for clarity; however, the scope is more affirmatively limited to a method for the specified seven steps.

Now compare and contrast two transatlantic claim counterparts for driverless car technology using reinforcement learning. First, European Patent No. 3,564,861, Vision-based sample-efficient reinforcement learning framework for autonomous driving, which is assigned to Sony.

“1. A method comprising: training a reinforcement learning controller for autonomous driving utilizing a vision model; and deploying the reinforcement learning controller for autonomous driving utilizing the vision model.”¹⁵³

The claim is concise at 44 total words, providing broad legal protection. Yet the quality is offset by preamble, which fails to state the method’s utility. The scope for this claim is extremely broad – at risk of being overly broad. Indeed, many researchers and companies are actively using, making, and selling driverless cars using reinforcement learning with computer vision.¹⁵⁴

¹⁵² European Patent 3129839A1, Controlling a target system (February 15, 2017). (Assigned to Siemens)

¹⁵³ European Patent 3564861A1, Vision-based sample-efficient reinforcement learning framework for autonomous driving (2019) (Assigned to Sony).

¹⁵⁴ See Noa Garnett, et. al., *Real-time Category-based and General Obstacle Detection for Autonomous Driving*, General Motors R&D, 198 (2017), http://openaccess.thecvf.com/content_ICCV_2017_workshops/papers/w3/Garnett_Real-Time_Category-Based_and_ICCV_2017_paper.pdf. See also Alex Kendall, et. al., *Learning to Drive in A Day*, 3 (2018) <https://arxiv.org/abs/1807.00412>. See also Ahmad El Sallab, et. al., *Deep Reinforcement Learning Framework for*

For example, United States Patent No. 10,061,316 Control policy learning and vehicle control method based on reinforcement learning without active exploration, is assigned to Toyota.

“1. A computer-implemented method for autonomously controlling a vehicle to perform a vehicle operation, the method comprising steps of: applying a passive actor-critic reinforcement learning method to passively-collected data relating to the vehicle operation, to adapt an existing control policy so as to enable control of the vehicle by the control policy so as to perform the vehicle operation with a minimum expected cumulative cost, the step of applying a passive actor-critic reinforcement learning method to passively-collected data including steps of: a) in a critic network, estimating a Z-value and an average cost under an optimal control policy using samples of the passively collected data; b) in an actor network operatively coupled to the critic network, revising the control policy using samples of the passively collected data, the estimated Z-value, and the estimated average cost under an optimal control policy from the critic network; and c) iteratively repeating steps (a)-(b) until the estimated average cost converges; and controlling the vehicle in accordance with the adapted control policy to perform the vehicle operation.”¹⁵⁵

The preamble is unnecessarily long, convoluting the invention’s subject matter. However, the transitional phrase comprising supports a broad reading. This is contrasted by the relatively high word count at 198, which narrows the claim because of the extended body.

Conventional wisdom says scope is a balancing of interest between broad claims,¹⁵⁶ and narrow claims.¹⁵⁷ This is a fallacy. Lemley explains, in patent litigation parties often successfully argue a claim means something completely different than the prosecuting attorney argued before the examiner.¹⁵⁸ A government product, the patent system is a game in which money wins.

Autonomous Driving, IS&T Electronic Imaging, *Autonomous Vehicles and Machines* 2017, AVM-023, 70, 71 (2017) <https://arxiv.org/pdf/1704.02532.pdf>.

¹⁵⁵ U.S. Patent No. 10,061,316 to Nishi, Control policy learning and vehicle control method based on reinforcement learning without active exploration (August 28, 2018). (Assigned to Toyota)

¹⁵⁶ Mark A. Lemley & Mark P. McKenna, *Scope*, 57 WM. & MARY L. REV. 2197, 2202 (2015).

¹⁵⁷ *Id.*

¹⁵⁸ Colleen Chien, Predicting Patent Litigation, 90 TEXAS L. REV. 283, 313-321 (2011), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1911579. (Providing empirical analysis relating to litigated and unlitigated patents.) See also Colleen Chien, Startups and Patent Trolls, 17 STAN. TECH. L. REV. 461 (2014), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2146251. See also Dan L. Burk & Mark A. Lemley, *Policy Levers in Patent Law*, 89 VA. L. REV. 1575, 1675 (2003).

Indeed, bureaucrats, agencies, and other governing organizations fail to perform basic tasks relating to patent information organization across decentralized information networks and silos.¹⁵⁹ As such, different bureaucrats decide doctrinal issues at different times through arbitrary decisions.¹⁶⁰ Nonetheless, the claim scope determines the monopolized market share.

IV. Property Strategy

One of the law's biggest secrets is that all inventions are patentable property.¹⁶¹ By definition an invention is a new, useful, and nonobvious machine or process.¹⁶² Not all ideas are inventions, but all inventions are property. And, reinforcement learning programs are valuable property. Consider the global market for AI technology now exceeds \$10.7 trillion.¹⁶³ Reinforcement learning systems drive the bleeding edge in this frontier. The way in which firms protect intellectual property is critical to valuation, growth and success.¹⁶⁴ Indeed, while most still think technology valuation is more art than science, more chance¹⁶⁵ than choice. Critically, technology is not randomly valuable. And, the way a technology is protected matters.

A. Patent Protections

¹⁵⁹ Veronica Root, *Coordinating Compliance Incentives*, 102 CORNELL L. REV. 1003, 1029 (2017), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2867048. (“As a result, the traditional concern regarding interagency coordination is with overlapping delegations of power.”)

¹⁶⁰ Eileen M. Herlihy, *The Ripple Effect of Seventh Amendment Decisions on the Development of Substantive Patent Law*, 27 SANTA CLARA COMPUTER & HIGH TECH. L.J. 333, 343 (2011).

¹⁶¹ USPTO Fee Schedule (2020), <https://www.uspto.gov/learning-and-resources/fees-and-payment/uspto-fee-schedule#Patent%20Fees>.

¹⁶² 35 U.S.C. § 101.

¹⁶³ Calculated as a share in the AI market. See United States Government Accountability Office, *Technology Assessment, Artificial Intelligence* (2018). (Stating the global impact for AI is \$10.7 Trillion.) See also McKinsey Global Institute, *Notes from the AI Frontier*, (2019). (Stating the global impact for AI is \$10.7 Trillion.)

¹⁶⁴ Anna Maria Baumgartner, *Antitrust Issues in Technology Transfer: A Comparative Legal Analysis of Patent Licenses in the EU and the U.S.*, TTLF Working Papers 145 (2013), <https://law.stanford.edu/publications/antitrust-issues-in-technology-transfer-a-comparative-legal-analysis-of-patent-licenses-in-the-eu-and-the-u-s/>. (“A significant convergence of EU and U.S. antitrust law can be observed when both approaches appraise licensing as potentially procompetitive.”)

¹⁶⁵ EDWARD O. THORPE, *A MAN FOR ALL MARKETS* 63-65 (2017).

To computer scientist, a patent is a document with invention data.¹⁶⁶ In this view, patents are made up with information bytes, binary logic compressed to macro-scales.¹⁶⁷ To the Economist, “Patents are like lotteries, in which there are a few prizes and a great many blanks.”¹⁶⁸ Speculating on price, the economists use market models based on random chance to value patents.¹⁶⁹ For most lawyers, patents are a strategic tool.¹⁷⁰ Patents have been awarded for inventions since the early 15th century.¹⁷¹ Despite progress since that time, public patent knowledge is deeply lacking due to systemic smoke and mirrors, complexities disguising more simple truth. The currency of kings, patents are limited monopolies for making, using, and selling inventions.

According to the highest authority in the land, patents are monopolies.¹⁷² The United States Patent and Trademark Office (“USPTO”) reviews applications to determine whether a

¹⁶⁶ Michael Buckland, *Information and Society* 21-23 (2017). (Discussing document definitions.)

¹⁶⁷ Mauritz Kop, *Machine Learning & EU Data Sharing Practices*, *Transatlantic Antitrust and IPR Developments*, Bimonthly Newsletter, Issue No. 1/2020, 7 (March 19, 2020). (“Data, or information, has a large number of legal dimensions.”)

¹⁶⁸ *The Economist*: Volume 9, Part 2, 811 (January 1, 1851). *See also* Edward O. Thorpe, *Blackjack Systems* (1975), <http://www.edwardothorp.com/wp-content/uploads/2016/11/BlackjackSystems.pdf>. *See also* Michael Gallagher II, *The Coupon Quandary: Restructuring Incentives in CAFA Coupon Settlements*, 91 *NOTRE DAME L. REV.* 2091, 2091 (2016). (Discussing attorneys fees in multi-million-dollar litigation.)

¹⁶⁹ *See* Albert Einstein, *Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen*, 322 *Annalen der Physik* 549 (1905). *See also* Robert Pitkethly, *The Value of Patents*, *Judge Institute Working Paper WP 21/97* (1997). *See also* Christopher A. Cotropia, *Describing Patents as Real Options*, 34 *J. Corp. L.* 1127 (2009). *See also* Mark A. Lemley, *Distinguishing Lost Profits from Reasonable Royalties*, 51 *WM. & MARY L. REV.* 655, 669 (2009).

¹⁷⁰ Shawn P. Miller, *Is There A Relationship Between Industry Concentration and Patent Activity?*, 3 (2009), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1531761.

¹⁷¹ Stefania Fusco, *Lessons from the Past: The Venetian Republic’s Tailoring of Patent Protection to the Characteristics of the Invention*, 17 *Nw. J. Tech. & Intell. Prop.* 101, 114 (2020). (“The first patent incorporating exclusionary rights that we know of was granted by the *Maggior Consiglio* in 1416 to *Ser Francisus Petri*, perhaps unsurprisingly a foreigner, for a device to full wool”)

¹⁷² U.S. Const. art. I, § 8, cl. 8. (Providing the constitutional basis for patents, “The Congress shall have the Power...To promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries.”) *See also* Andrew Beckerman-Rodau, *The Problem with Intellectual Property Rights: Subject Matter Expansion*, 13 *Yale J. L. & Tech.* 36, 55 (2010-2011). (The USPTO’s granting of patent rights provides typical property rights, including the right of the patent owner to exclude others from making, using, offering for sale, or selling the invention throughout the United States or importing the invention into the United States.) *See also* JOHN PALFREY, *INTELLECTUAL PROPERTY STRATEGY* 55 (MIT Press 2012). *See also* Stephen Yelderman, *The Value of Accuracy in The Patent System*, 84 *U. CHI. L. REV.* 1217, 1270 (2017).

claimed invention is: (1) statutory subject matter,¹⁷³ (2) useful,¹⁷⁴ (3) novel,¹⁷⁵ (4) nonobvious,¹⁷⁶ and (5) sufficiently described.¹⁷⁷ Critically under U.S. Law, new, useful, and nonobvious reinforcement learning software and processes are patentable subject matter.¹⁷⁸ Upon issuance, patents confer the exclusive right to use and profit from an invention to the holder, a government backed monopoly.¹⁷⁹ A U.S. Patent's term extends twenty years from the date the USPTO grants the patent.¹⁸⁰ The total procedural fee for acquiring a U.S. Patent is \$2,420.00, which is discounted to \$1,285.00 for small entities, and to \$605.00 for micro entities.¹⁸¹

¹⁷³ 35 U.S.C. § 101. (The first element of the statutory requirements, statutory subject matter, includes any new process, machine, manufacture, or composition of matter, or any new and useful improvement thereof.)

¹⁷⁴ Generally, a good test for utility is to identify a problem the invention solves.

¹⁷⁵ Professor Elona Marku at the University of Cagliari in Italy is developing objective measures for novelty. In fact, they have developed a quality formalism for measuring patent originality, which may be modified to measure novelty. *See* Elona Marku, et al., *Quantity at expense of quality? Measuring the effects of "successful" M&A on innovation performance* 8 (2020).

¹⁷⁶ Jeanne C. Fromer, *The Layers of Obviousness in Patent Law*, 22 *HARV. J. OF L. & TECH.*, 75 (2008). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1119723. ("The nonobviousness doctrine seeks to ensure that patents are granted only for technologically significant advances to foster the patent system's goal of stimulating useful innovation.")

¹⁷⁷ 35 U.S.C. § 112. *See also* Sarah Murphy, *Heads I Win, Tails You Lose: The "Expense" of a De Novo Review of USPTO Decisions*, 60 *B.C.L. Rev.* II.-197, II.-197 (2019), <https://lawdigitalcommons.bc.edu/bclr/vol60/iss9/15>. ("The United States Patent and Trademark Office (the "USPTO") may deny patent applications and trademark registrations to applicants who do not meet the necessary requirements.")

¹⁷⁸ Julie E. Cohen and Mark A. Lemley, *Patent Scope and Innovation in the Software Industry*, 89 *Cal. L. Rev.* 1, 8 (2001). ("Today, it seems fairly settled that software-related inventions fall within the class of innovations described in section 101 of the Patent Act as eligible for patent protection.") *See also* Burk, Dan L. and Lemley, Mark A., *Quantum Patent Mechanics*, *Lewis & Clark Law Review*, Vol. 9, p. 29, 30 (2005), <https://ssrn.com/abstract=628224>. ("Determining that meaning necessarily requires the judge to break the text of a claim into discrete "elements" or units of text corresponding to the elements or units that comprise the claimed invention-essentially, organizing the language of the claims into "chunks" or "quanta" of text.")

¹⁷⁹ U.S. Const. art. I, § 8, cl. 8. (Providing the constitutional basis for patents, "The Congress shall have the Power... To promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries.") *See also* Stephen Yelderman, *The Value of Accuracy in The Patent System*, 84 *U. CHI. L. REV.* 1217, 1270 (2017). *See also* Bryce C. Pilz, *Student Intellectual Property Issues on the Entrepreneurial Campus*, 2 *MICH. J. PRIVATE EQUITY & VENTURE CAP. L.* 1, 16 (2012).

¹⁸⁰ 35 U.S.C. 154 (2020).

¹⁸¹ USPTO Fee Schedule (2020), <https://www.uspto.gov/learning-and-resources/fees-and-payment/uspto-fee-schedule#Patent%20Fees>. *See* 37 CFR 1.16(a); 37 CFR 1.16(k); 37 CFR 1.16(o); 37 CFR 1.18(a)(1). *See also* Shawn P. Miller, *What's The Connection Between Repeat Litigation And Patent Quality? A (Partial) Defense of The Most Litigated Patents*, 16 *Stan. Tech. L. Rev.* 313, 317 (2013). ("Thus, over 98 percent of granted patents never generate direct litigation costs.")

The European Patent Office grants patents for inventions, which are new, novel, and industrially applicable technologies or innovations.¹⁸² Similar to the United States patent application, the European patent application¹⁸³ requires: (1) an abstract, (2) a detailed description, (3) claims, and (4) drawings.¹⁸⁴ Also, similar to the United States inventions utilizing reinforcement learning software are patentable subject matter.¹⁸⁵ Distinctly, European Patents provide territorial protection and European Patents¹⁸⁶ guarantee exclusive rights in thirty-eight member states.¹⁸⁷ The total procedural cost for a European Patent is about €4,135.00.¹⁸⁸

¹⁸² The European Patent Convention, Article 52, Patentable inventions (2020), <https://www.epo.org/law-practice/legal-texts/html/epc/2016/e/ar52.html>. (“European patents shall be granted for any inventions, in all fields of technology, provided that they are new, involve an inventive step and are susceptible of industrial application.”) *See also* European IPR Helpdesk, IPR Chart (2020). (“Patents protect technical inventions, which are new (do not form part of the state of the art), industrially applicable and involve an inventive step (not obvious to a person skilled in the art).”) *See also* European Patent Office, Guidelines for Examination in the European Patent Office, General Part 1, 15 (2019). (“In accordance with Art. 10(2)(a) of the European Patent Convention (EPC), the President of the European Patent Office (EPO) had adopted, effective as at 1 June 1978, the Guidelines for Examination in the European Patent Office.”)

¹⁸³ European Patent Office, Guidelines for Examination in the European Patent Office, Part A, 41 (2019). (“European patent applications may be filed by delivery by hand or by postal services at the EPO’s filing offices in Munich, The Hague or Berlin.”)

¹⁸⁴ European IPR Helpdesk, IPR Chart (2020). (“An EP application must contain: (1) a request for the grant of an EP, (2) a description of the invention, where all details of the invention together with references to prior art are given, (3) one or more claims, which define the scope of protection, (4) drawings, if any, and (5) an abstract.”)

¹⁸⁵ The European Patent Convention, Article 52, Patentable inventions (2020), <https://www.epo.org/law-practice/legal-texts/html/epc/2016/e/ar52.html>.

¹⁸⁶ Convention on the Grant of European Patents, 8 (2000). (“Patents granted under this Convention shall be called European patents.”) *See also* Siegfried Fina & Anna Maria Baumgartner, A Comparative Antitrust Analysis of Exclusivity Clauses in Patent Licenses Under Article 101 TFEU and Section 1 Sherman Antitrust Act, TTLF Working Papers, 3 (2012). (“Article 101 TFEU contains the general antitrust ban in EU law by prohibiting all 1) agreements between parties, 2) decisions by associations of undertakings, and 3) concerted practices, which may affect trade between Member States and which have, as their object or effect, the prevention, restriction or distortion of competition within the internal marketplace.”)

¹⁸⁷ Convention on the Grant of European Patents, 8 (2000). (“Currently the 38 Contracting States are: AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LI, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR.”) *See also* European IPR Helpdesk, IPR Chart (2020). (“Patent rights are territorial by nature. This means that the protection is only valid in countries in which that patent is granted.”) *See also* Robert Pitkethly, The European Patent System: Implementing Patent Law Harmonisation, 8 (1999). (“If one ignores the post grant life of a European Patent Application as a set of national patents and distinguishes the European Patent Convention from National Patent Systems in terms of the extra functions and costs the unitary European layer adds to the patent application process; it can be seen that the key benefits reside primarily in cost savings and efficiency but also in issues involving the uniformity and certainty of protection.”)

¹⁸⁸ Filing a patent before the European Patent office is €125.00 Euros; the search fee for patent applications is €1,350.00; the examination fee is €1,700.00 Euros;¹⁸⁸ the publishing fee for grant and publish is €960.00 Euros. European Patent Office, Online Fee Payment – Schedule of Fees (2020).

Consider the follow example claim for a reinforcement learning application in nanotechnology for healthcare.¹⁸⁹ The person may claim the following:¹⁹⁰

A method for autonomous thyroid hormonal control,¹⁹¹ the method consisting of a synthetic module¹⁹² delivering the thyroid hormone Thyroxine (T₄),¹⁹³ to a patient's bloodstream on a timed basis, using a reinforcement learning agent,¹⁹⁴ optimizing patient Thyroxine (T₄) levels at a consistent homeostasis.

This claim is narrowed by the transitional phrase consisting, but otherwise provides a sustainable process for controlling thyroid hormonal regulation in the human body using reinforcement learning. Reinforcement learning's integration with nanotechnology presents an opportunistic plethora. As Ouellette explains, "Nanotechnology is thus a useful counterpoint both to the growing number of case studies on how innovation can flourish without intellectual property ("IP"), and to the myth of an independent private sector that produces breakthrough innovations without government intervention."¹⁹⁵

Similarities between the United States and European patent systems are patentable subject matter and term.¹⁹⁶ Two differences are EPO patents are slightly more expensive and protect in 38 countries, compared to the USPTO patents which protect only in the United

¹⁸⁹ Lisa Larrimore Ouellette, *Nanotechnology and Innovation Policy* 29 HARV. J. L. & TECH. 33, 34 (2015). ("Nanotechnology is the engineering of matter at scales less than about one hundred nanometers (one ten-millionth of a meter)"). See also Simone Schuerle, et al., *Three-Dimensional Magnetic Manipulation of Micro- and Nanostructures for Applications in Life Sciences* (2013), <https://ieeexplore.ieee.org/document/6392417>.

¹⁹⁰ Mark A. Lemley, Samantha Zyontz, *Does Alice Target Patent Trolls?*, 4 (2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3561252. ("In the most recent decision, *Alice*, the Court confirmed a two-step test for determining whether a claim was patentable: (1) is the claim "directed to" an abstract idea, law of nature, or other excluded subject matter? and (2) if so, does the claim include an inventive step beyond merely the claimed abstract idea or natural phenomenon?")

¹⁹¹ The term "Autonomous Metabolic Control" refers to regulating the human metabolism automatically.

¹⁹² A "Synthetic Module" is a sterilized enclosed capsule.

¹⁹³ The term "Thyroxine (T₄)" is a thyroid hormone.

¹⁹⁴ A "Reinforcement Learning Agent" is a program capable of learning, taking actions, and achieving goals.

¹⁹⁵ Lisa Larrimore Ouellette, *Nanotechnology and Innovation Policy* 29 HARV. J. L. & TECH. 33, 34 (2015).

¹⁹⁶ European IPR Helpdesk, *IPR Chart* (2020). ("The European Patent is valid for 20 years from the EP filing date as long as the yearly maintenance fees are paid within this period in designated countries.")

States.¹⁹⁷ Some argue the patent system as a whole is moving toward more uniformity,¹⁹⁸ but skeptics suggest litigation costs deter firms from patent investment altogether.¹⁹⁹ Maybe patents are a dying art.²⁰⁰ Obviously, secrets are more fun.

B. Software Secrets

Thiel argues proprietary technology is the most substantive advantage a company can have because the secret nature makes the product difficult to replicate.²⁰¹ Consider instead the only variable for reverse engineering any technology is time. Unlike patent protections, which take years, bringing more liability than revenue, trade secret protections are immediate and free. In fact, trade secret protections simply require the firm take reasonable steps to prevent disclosure and protect the underlying invention.²⁰² Hrdy and Lemley argue trade secrets give rights to those who have yet had an opportunity to acquire the means to use for building or using technology.²⁰³

¹⁹⁷ Convention on the Grant of European Patents, 8 (2000). *See also* USPTO Fee Schedule (2020), <https://www.uspto.gov/learning-and-resources/fees-and-payment/uspto-fee-schedule#Patent%20Fees>.

¹⁹⁸ Lisa Larrimore Ouellette, Patent Experimentalism, 101 Va. L. R. 65, 67 (2015), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2294774. (“The dominant push in patent law, however, has been toward uniformity.”)

¹⁹⁹ Shawn P. Miller, *What's The Connection Between Repeat Litigation And Patent Quality? A (Partial) Defense of The Most Litigated Patents*, 16 Stan. Tech. L. Rev. 313, 314 (2013). (“Skeptics of the U.S. patent system’s recent effectiveness in promoting innovation fear the impact of increased costs imposed on firms that bring new products to market because of the explosion in patent litigation during the 1990s.”)

²⁰⁰ For example, most people think the only reason pharmaceutical companies are willing to pay the high costs associated FDA approval is they can charge high prices for patented drugs. But the truth is if FDA approval was a public good, then the costs would be zero.

²⁰¹ PETER THIEL, *ZERO TO ONE* 48 (2014)

²⁰² One way in which this may be accomplished is by simply marking the document describing the invention itself. *See* Trade Secret, Black’s Law Dictionary (10th ed. 2014). *See also* Mark A. Lemley, The Surprising Virtues of Treating Trade Secrets as IP Rights, 61 Stan. L. Rev. 311, 329 (2008). *See also* Camilla A. Hrdy, Mark A. Lemley, Abandoning Trade Secrets, 1 (2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3534322.

²⁰³ Camilla A. Hrdy, Mark A. Lemley, Abandoning Trade Secrets, 3 (2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3534322. (Citing UTSA § 1, cmt.) (“Instead, today’s independent economic value criterion, which is derived from the common law concept of “competitive advantage,” envisions various ways to derive economic value from information, ranging from selling or licensing information to others to suppressing early stage research related to a product; and it gives rights to those who have “not yet had an opportunity or acquired the means to put a trade secret to use[.]””)

In *Ruckelshaus v. Monsanto Co.*, the United Supreme Court held trade secrets are a form of property, which may be protected by the Constitution.²⁰⁴ The Court reasoned trade secrets possess many characteristics of tangible property, for example they are assignable.²⁰⁵ New law²⁰⁶ developed to allow large firms to sue former employees for trade secret misappropriation. However, the Defend Trade Secrets Act (DTSA) has more important consequences for property strategy.

Recognizing trade secrets under federal law, Congress aggrandized the firm's ability to financially protect proprietary technology.²⁰⁷ According to the DTSA:

“An owner of a trade secret that is misappropriated may bring a civil action under this subsection if the trade secret is related to a product or service used in, or intended for use in, interstate or foreign commerce.”²⁰⁸

Yelderman argues, the added protections at the federal level make firms more likely to pursue trade secret protections as opposed to traditional patent protections.²⁰⁹ Yet, despite the DTSA's added security, protecting confidential information in an arena with contrasting interests is a complex task.²¹⁰

According to the European Union Intellectual Property Office (EUIPO), trade secrets are critically important to the growth competitiveness and innovative performance.²¹¹ Generally, the

²⁰⁴ *Ruckelshaus v. Monsanto Co.*, 467 U.S. 986 (1984). (“We therefore hold that to the extent that Monsanto has an interest in its health, safety, and environmental data cognizable as a trade-secret property right under Missouri law, that property right is protected by the Takings Clause of the Fifth Amendment.”)

²⁰⁵ Mark A. Lemley, *The Surprising Virtues of Treating Trade Secrets as IP Rights*, 61 *Stan. L. Rev.* 311, 329 (2008).

²⁰⁶ Jeanne C. Fromer, *Machines as the New Oompa-Loompas: Trade Secrecy, the Cloud, Machine Learning, and Automation*, N.Y.U. L.R., 706, 709 (2019), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3359746. (“Until recently, U.S. trade secret law was principally state-based.”)

²⁰⁷ Mark A. Lemley, *The Surprising Virtues of Treating Trade Secrets as IP Rights*, 61 *STAN. L. REV.* 311, 329 (2008).

²⁰⁸ Defend Trade Secrets Act of 2016 (May 11, 2016).

²⁰⁹ Stephen Yelderman, *The Value of Accuracy in The Patent System*, 84 *U. CHI. L. REV.* 1217, 1264 (2017).

²¹⁰ Suellen Lowry, *Inevitable Disclosure Trade Secret Disputes: Dissolutions of Concurrent Property Interests*, 40 *STAN. L. REV.* 519, 519 (1998).

²¹¹ European Union Intellectual Property Office, *The Baseline of Trade Secrets Litigation in the EU Member States*, 5 (2018), doi: 10.2814/19869. (“Trade secrets and confidential business information are critically important to the

EUPO bases this analysis on the idea that the current business landscape values intellectual capital as a competitive advantage.²¹² According to EUPO: “Where innovation does not fulfil the requirements of patentability, trade secrets become a pivotal tool for companies to protect their business knowledge.”²¹³ But consider a different truth: there is no competition in modern markets, innovation is dead, and the only intellectual capital is cash.

All proprietary technology can be reverse engineered with reinforcement learning.²¹⁴ For example, rocket technology was previously thought to be a highly classified and protected proprietary technology.²¹⁵ According to SpaceX Founder & CEO Elon Musk, “our primary long-term competition is China – if we published patents, it would be farcical because the Chinese would just use them as a recipe book.”²¹⁶ But the internet alone is a fine recipe book. For example, consider the following claim:

A method for landing rockets, the method consisting of a rocket optimizing safe landing metrics, controlling attitude, pitch, roll, and yaw with thrust vector commands, manipulating thruster output, using a reinforcement learning system: the system

growth, competitiveness and innovative performance of European businesses.”) *See also* Siegfried Fina and Gabriel M. Lentner, The Scope of the EU's Investment Competence after Lisbon, 14 Santa Clara J. Int'l L. 419, 420 (2016), <http://digitalcommons.law.scu.edu/scujil/vol14/iss2/2>. (“The increased importance of global investment flows means that rules on investment promotion and protection are vital in stimulating trade relations and have a positive influence on the quality and quantity of investments.”)

²¹² European Union Intellectual Property Office, The Baseline of Trade Secrets Litigation in the EU Member States, 10 (2018), doi: 10.2814/19869. (“In the current business landscape, companies and research organisations invest in generating and acquiring intellectual capital to obtain a competitive advantage and foster their innovation-related performances.”)

²¹³ European Union Intellectual Property Office, The Baseline of Trade Secrets Litigation in the EU Member States, 5 (2018), doi: 10.2814/19869. (“Where innovation does not fulfil the requirements of patentability, trade secrets become a pivotal tool for companies to protect their business knowledge.”)

²¹⁴ Julie E. Cohen and Mark A. Lemley, Patent Scope and Innovation in the Software Industry, 89 Cal. L. Rev. 1, 18 (2001). (“Because reverse engineering is costly, this legal rule does not foreclose the possibility of a licensing arrangement. But it does prevent a potential licensor from refusing to deal at all, and it imposes a natural upper limit—the cost of reverse engineering—on what a licensee will be willing to pay.”) *See also* Olia Kanevskaia, The law and practice of global ICT standardization, 116 (March 31, 2020). (“But despite the challenges “open standards,” such as inclusion of proprietary solutions into Internet specifications and lack of sufficient governmental recognition, this concept, as well as the OpenStand principles, cannot be ignored in the context of modern standardization.”)

²¹⁵ *Space Exploration Technologies v. Blue Origin*, USPTO Trial and Appeal Board, Judge Carl M. DeFranco, et. al. (2015). *See also* U.S. Patent No. 8,678,321, to Bezos, et. al. Sea landing of space launch vehicles and associated systems and methods (March 25, 2014).

²¹⁶ Anderson, *supra* note 40.

consisting of a convolutional neural network predicting trajectory position, feeding environmental data to an agent, taking actions and landing performance.

The claim scope is limited by the transitional phrase “consisting” which is present twice. From a financial perspective, utilizing reinforcement learning for rocket technology provides opportunity for great space exploration and financial return.²¹⁷

Compared to trade secrets, patent rights have more lucid boundaries due to public disclosure.²¹⁸ But, trade secrets cost significantly less and both patents and trade secrets may be used successfully as leverage in negotiations. In theory, patents benefit society more broadly than trade secrets since patents are in the public domain.²¹⁹ But all code for reinforcement learning should also include copyright protections, which like trade secrets are free.

C. Code Copyrights

Distinct from patents and trade secrets, which protect inventions and proprietary technology, copyright law does not protect novel machines or processes.²²⁰ Koski considers the critical role software copyrights play in the licensing context.²²¹ Lemley and Casey recently analyzed copyright law specifically as it pertains to machine learning software.²²² They argue

²¹⁷ Brian Bozzo, *Not Because Its Easy: Exploring National Incentives for Commercial Space Exploration through a Geopolitical Lens*, 11 Drexel L. Rev. 595, 598 (2019). (“By 1973, taxpayers had paid the 2018 equivalent of about \$200 billion in exchange for six successful moon landings and proof of “the superiority of a democracy.”)(“On November 11, 2018, Rocket Lab launched its first fully commercial mission to low-Earth orbit on the “It’s Business Time” rocket. The launch cost the company a mere \$5 million.”)

²¹⁸ Colleen Chien, *Software Patents as a Currency, Not Tax, on Innovation*, 31 BERKELEY TECH. L.J. 1669, 1681 (2017), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2870813. (The boundaries of patent rights are also more readily ascertainable than trade secrets, defining the duration of the right and the scope of the claims so that the parties do not have to do so.)

²¹⁹ Max Stul Oppenheimer, *Patents 101: Patentable Subject Matter and Separation of Powers*, 15 Vand. J. Ent. & Tech. L. 1, 7 (2012).

²²⁰ 17 U.S.C. §102. “(b) In no case does copyright protection for an original work of authorship extend to any idea, procedure, process, system, method of operation, concept, principle, or discovery, regardless of the form in which it is described, explained, illustrated, or embodied in such work.”

²²¹ Heli Koski, *OSS Production and Licensing Strategies of Software Firms* 111, 113 (2005). (“Software producers have a variety of options to distribute their products.”)

²²² Mark A. Lemley, Bryan Casey, *Fair Learning*, 5 (2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3528447. (“After decades of allowing – or even just plain ignoring – machine copying, copyright owners and courts have begun to loudly and visibly push back against the copyright system’s permissive attitude towards machine copying.”)

machine learning “systems should generally be able to use databases for training, whether or not the contents of that database are copyrighted.”²²³

In the United States, computer programs have been copyrightable since 1976.²²⁴ Consider 17 U.S.C. §102, which defines subject matter which may be afforded copyright protection:

“(a) Copyright protection subsists, in accordance with this title, in original works of authorship²²⁵ fixed in any tangible medium of expression, now known or later developed, from which they can be perceived, reproduced, or otherwise communicated, either directly or with the aid of a machine or device.”²²⁶

One question is whether copyright law protects against reverse engineering. Oman argues the answer is yes, “[F]or purposes of understanding copyright protection for computer software is that a literary work can be both functional and expressive.”²²⁷ However, according to Cohen and Lemley, the answer is no, “[V]irtually every court to consider the issue has concluded that there is a right to reverse engineer a copyrighted program for at least some purposes.”²²⁸

²²³ Mark A. Lemley, Bryan Casey, *Fair Learning*, 7 (2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3528447. (“In this Article, we argue that ML systems should generally be able to use databases for training, whether or not the contents of that database are copyrighted.”)

²²⁴ Lisa C. Green, Copyright Protection and Computer Programs: Identifying Creative Expression in a Computer Program’s Nonliteral Elements, 3 *Fordham Intell. Prop. Media & Ent. L.J.* 89, 89 (1992). See also Ralph Oman, Computer Software as Copyrightable Subject Matter: Oracle v. Google, Legislative Intent, and The Scope of Rights in Digital Works 31 *HARV. J. L. & T.* 638, 641 (2018). (“In 1980, as the world sat on the brink of the digital age, Congress amended the Copyright Act to provide computer programs the same copyright protection as all other literary works.”)

²²⁵ 17 U.S.C. §102. “Works of authorship include the following categories: (1) literary works; (2) musical works, including any accompanying words; (3) dramatic works, including any accompanying music; (4) pantomimes and choreographic works; (5) pictorial, graphic, and sculptural works; (6) motion pictures and other audiovisual works; (7) sound recordings; and (8) architectural works.” See also 17 U.S.C. § 101. Definitions (2020). ““Literary works” are works, other than audiovisual works, expressed in words, numbers, or other verbal or numerical symbols or indicia, regardless of the nature of the material objects, such as books, periodicals, manuscripts, phonorecords, film, tapes, disks, or cards, in which they are embodied.”

²²⁶ 17 U.S.C. §102.

²²⁷ Ralph Oman, Computer Software as Copyrightable Subject Matter: Oracle v. Google, Legislative Intent, and The Scope of Rights in Digital Works 31 *HARV. J. L. & T.* 638, 644 (2018). (“Perhaps the most important principle for purposes of understanding copyright protection for computer software is that a literary work can be both functional and expressive.”)

²²⁸ Julie E. Cohen and Mark A. Lemley, Patent Scope and Innovation in the Software Industry, 89 *Cal. L. Rev.* 1, 17 (2001). (“While there is no express statutory provision in the copyright laws permitting reverse engineering, virtually every court to consider the issue has concluded that there is a right to reverse engineer a copyrighted program for at least some purposes.”)

Providing copyright notice pursuant to 17 U.S.C. §401 simply requires: (1) the word “Copyright”; (2) the year of the work’s first publication; and (3) the owner’s name.²²⁹ Given rising complexity in code and corresponding copyright protections, problems persist for programmers, firms, and universities.²³⁰ For example, the Copyright Act is generally poorly written, leading many to question Congressional intent in interpretation.²³¹ In turn, more words and complex statutes yield less clarity and more deference to courts.

In Europe, copyrights derive from a set of international principles, norms, and treaties,²³² including EU directives, national laws, and case law.²³³ Computer code has been copyrightable subject matter under EU law for nearly thirty years pursuant to the EU’s Software Directive.²³⁴ According to Samuelson, “The Directive could not be clearer in its conception of computer programs as copyright subject matter.”²³⁵

²²⁹ 17 U.S.C. §401; Notice of copyright: Visually perceptible copies (2020).

²³⁰ See Dr. Roland Vogl, et al., Addressing the Copyright Law Barrier in Higher Education – Access-to-Clean-Content Technology in the 21st Century, Stanford CodeX, 7 (2012),

<https://law.stanford.edu/publications/addressing-the-copyright-law-barrier-in-higher-education-access-to-clean-content-technology-in-the-21st-century/>. (“The U.S. lawsuit *Cambridge University Press v. Patton* exposes university vulnerabilities regarding fair use and data practices. In that case, the defendant “Georgia State University’s policies and training of faculty regarding copyright issues did not protect it from liability.”) See also *University Press v. Becker*, No. 1:08-cv-01425- ODE, 2012 BL 119127 (N.D. Ga. May 11, 2012).

²³¹ Ralph Oman, Computer Software as Copyrightable Subject Matter: Oracle v. Google, Legislative Intent, and The Scope of Rights in Digital Works 31 HARV. J. L. & T. 638, 640 (2018). (“But when Congress made clear that computer programs would, in fact, be copyrightable, it effectively imported several centuries’ worth of well-studied (if often misunderstood) legal doctrines to bear on the questions that would inevitably follow from its decision to categorize software as a “literary work,” protectable like any other.”)

²³² Mauritz Kop, AI & Intellectual Property: Towards an Articulated Public Domain 28 Tex. I.P. L. J., 1, 5 (2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3409715. (“The main IP Treaties are the Berne Convention, the Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS) and the WIPO Copyright Treaty (WTC), to which almost all countries of the world are members. This multilevel framework means that, at a national level in both USA and EU, there is an obligation to comply with international treaties.”)

²³³ Mauritz Kop, AI & Intellectual Property: Towards an Articulated Public Domain 28 Tex. I.P. L. J., 1, 5 (2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3409715. (“Copyright consists of a set of international principles, including rationales and justifications, and a set of norms, (laid down in) Treaties, EU Directives, national laws and case law.”)

²³⁴ Pamela Samuelson, *Does Copyright Protection Under the EU Software Directive Extend to Computer Program Behaviour, Languages and Interfaces?*, 1 (2011). (“Twenty years ago, European software copyright law was successfully harmonised by the EU’s 1991 Software Directive.”)

²³⁵ Pamela Samuelson, *Does Copyright Protection Under the EU Software Directive Extend to Computer Program Behaviour, Languages and Interfaces?*, 3 (2011).

The European Parliament explains, national copyright systems follow two different legal traditions: civil law and common law.²³⁶ Further, the European Parliament contends:

“[A]lthough copyright law in the European Union remains essentially national law, national rules are gradually converging by means of alignment with international treaties and Union legislation, which harmonise the various rights of authors, performers, producers and broadcasters.”²³⁷

Yet, Kanevskaia argues the EU is lacking in ruling on “access to copyrighted standards referenced in regulation.”²³⁸ Moreover, Samuelson echoes this concern discussing the delicacy in balancing interests defining copyright scope for computer programs,²³⁹ which would include reinforcement learning programs.

For reinforcement learning programs, a common copyright protection ascribed is the Apache License.²⁴⁰ For example, the software code for DeepMind’s deep Q-network on GitHub is marked with the following:

```
“# Copyright 2020 DeepMind Technologies Limited. All Rights Reserved.  
#
```

²³⁶ European Parliament, Copyright Law in the EU, 3 (June 2018). (“National copyright systems follow two different legal traditions: civil law in continental Europe and common law in the United Kingdom, Ireland, Malta and Cyprus.”)

²³⁷ European Parliament, Copyright Law in the EU, 1 (June 2018). (“However, although copyright law in the European Union remains essentially national law, national rules are gradually converging by means of alignment with international treaties and Union legislation, which harmonise the various rights of authors, performers, producers and broadcasters.”)

²³⁸ Olya Kanevskaia, *The law and practice of global ICT standardization*, 71 (March 31, 2020). (“Fortunately, unlike the EU, the US Courts have on a number of occasions ruled on the issue of access to copyrighted standards referenced in regulation.”) *See also* Eleonora Rosati, *Closed subject-matter systems are no longer compatible with EU copyright*, 6 (2014), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2468104. (“From the case law mentioned above, it appears that subject-matter categorisation under EU copyright should be intended as open-ended.”)

²³⁹ Pamela Samuelson, *Does Copyright Protection Under the EU Software Directive Extend to Computer Program Behaviour, Languages and Interfaces?*, 1 (2011). (“In defining the scope of copyright protection for computer programs for the EU, the drafters of the Directive sought to strike a careful balance between, on the one hand, providing appropriate copyright protection to computer programs in order to stimulate investments in new software development, and, on the other hand, enabling second comers to engage in independent development of software capable of fully interoperating with other programs.”)

²⁴⁰ Apache License, Version 2.0 (January 2004), <http://www.apache.org/licenses/>. (“2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.”)

```

# Licensed under the Apache License, Version 2.0 (the "License");
# you may not use this file except in compliance with the License.
# You may obtain a copy of the License at
#
# http://www.apache.org/licenses/LICENSE-2.0
#
# Unless required by applicable law or agreed to in writing, software
# distributed under the License is distributed on an "AS IS" BASIS,
# WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or
# implied.
# See the License for the specific language governing permissions and
# limitations under the License.”241

```

The general effect for open-sourcing license is to grant a free license, with a limited liability warranty. This methodology for open sourcing software is often advantageous for firms to drive licensing revenue because it introduces a limited software version for free, while retaining more important rights.²⁴² However, defining the protectable limits for this code are inherently complex considering overlap between copyrights and patent rights with DeepMind parent, Google.²⁴³

An opaque future awaits subject matter protections in both the United States and European Union.²⁴⁴ According to Vogl, one problem with copyright law is “publishers are losing

²⁴¹ Deepmind, deepmind-research (2020). https://github.com/deepmind/deepmind-research/blob/master/option_keyboard/gpe_gpi_experiments/run_dqn_fig5.py. See also APACHE LICENSE, VERSION 2.0 (2004), <http://www.apache.org/licenses/LICENSE-2.0>.

²⁴² Heli Koski, OSS Production and Licensing Strategies of Software Firms 111, 117 (2005). (“A firm’s business model and particularly the importance of license revenues are also likely to play an important role in defining the firm’s product and license type decisions.”) A similar strategy is used for Blockchain technologies. See Igor Linkov, Emily Wells, et al., Blockchain Benefits and Risks (2018), https://www.researchgate.net/profile/Igor_Linkov/publication/325385235_Blockchain_Benefits_and_Risks/links/5df6b251a6fdcc2837245f1e/Blockchain-Benefits-and-Risks.pdf. (“Blockchain methods—formally referred to as distributed ledger technologies—enable live, interconnected transaction records with data access and disclosure to all members of a network.”)

²⁴³ Ralph Oman, *Computer Software as Copyrightable Subject Matter: Oracle v. Google, Legislative Intent, and The Scope of Rights in Digital Works*, 31 Harv. J. L. & T. 639, 646 (2018). (“Nevertheless, identifying the line where protectable expression ends and unprotectable function begins requires defining both the work for which protection is sought, as well as the function it performs.”) See also Volodymyr Mnih, Koray Kavukcuoglu, Methods and Apparatus for Reinforcement Learning, U.S. Patent No. 9,679,258 B2 (2017) (<https://patents.google.com/patent/US9679258B2/en>). See also TensorFlow, GitHub, DQN (2020). https://github.com/tensorflow/agents/tree/master/tf_agents/agents/dqn (Code for DQN from TensorFlow under an Apache license.)

²⁴⁴ Olia Kanevskaia, The law and practice of global ICT standardization, 71 (March 31, 2020). (“As it is the case in the EU, copyrights of US standards referenced in law is increasingly becoming a topic of discussion.”)

track of where and how their content is being used, so they have little or no data to rely on to properly price content or know how or where to best create reliable revenue streams.”²⁴⁵ And, Lemley and Casey explain another specific issue for machine learning copyrights is whether AI technologies can process copyrighted material for training purposes.²⁴⁶ Indeed, copyright law is the wild card for reinforcement learning property strategy because the bounds are most blurred.

Conclusion

New technologies often result from a technical convergence including various techniques, devices, and machines.²⁴⁷ For example, personal computers and mobile phones converged, forming the smart phone.²⁴⁸ Elona Marku explains,²⁴⁹ convergence is “the merging or overlapping of different fields of technology as a result of scientific and technological progress.”²⁵⁰ Reinforcement learning is important because it is a scalable method for general

²⁴⁵ Dr. Roland Vogl, et al., Addressing the Copyright Law Barrier in Higher Education – Access-to-Clean-Content Technology in the 21st Century, Stanford CodeX, 2 (2012), <https://law.stanford.edu/publications/addressing-the-copyright-law-barrier-in-higher-education-access-to-clean-content-technology-in-the-21st-century/>. (“A further larger problem is that publishers are losing track of where and how their content is being used, so they have little or no data to rely on to properly price content, or know how or where to best create reliable revenue streams.”)

²⁴⁶ Mark A. Lemley, Bryan Casey, *Fair Learning*, 4 (2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3528447. (“Creating a training set of millions of examples almost always requires, first, copying many more millions of images, videos, audio, or text-based works. Those works are almost all copyrighted.”)

²⁴⁷ PAUL E. CERUZZI, *COMPUTING: A CONCISE HISTORY* 74-76 (2012).

²⁴⁸ Interestingly, in 2009 Nokia and Samsung paid a small semiconductor firm in King of Prussia, Pennsylvania called InterDigital a combined \$653 million over a portfolio of patents for smart phone technology. See JOHN PALFREY, *INTELLECTUAL PROPERTY STRATEGY* 18 (MIT Press 2012). See also *In Matter of Arbitration Between InterDigital Communications Corp. and Samsung...*, 528 F.Supp.2d 340 (2007). See also *InterDigital Communications Corp. v. Nokia Corp.*, 407 F.Supp.2d 522 (2005).

²⁴⁹ Elona Marku, et al., Mapping Innovation in the Digital Transformation Era: The Role of Technology Convergence, 165 (2019), https://www.researchgate.net/publication/329874675_Mapping_Innovation_in_the_Digital_Transformation_Era_The_Role_of_Technology_Convergence. (“This is consistent with the technological convergence paradigm, two or more technologies move together in the technological space, overlapping or merging with each other while generating new innovations.”)

²⁵⁰ Elona Marku, et al., Mapping Innovation in the Digital Transformation Era: The Role of Technology Convergence, 163 (2019), https://www.researchgate.net/publication/329874675_Mapping_Innovation_in_the_Digital_Transformation_Era_The_Role_of_Technology_Convergence. (“Consistent with this stream of research, we conceive technological convergence as the merging or overlapping of different fields of technology as a result of scientific and technological progress.”)

intelligence. As reinforcement learning models continue to converge with deep learning software, AI programs continue outperforming humans in virtually all tasks.

Convergence toward more open and transparent models may reflect in the property strategy for reinforcement learning systems as well. From a transatlantic perspective, open strategies make sense because the world's information is entirely decentralized across the internet, dismantling notions of truly proprietary or classified information. Further, open-source strategies allow firms to profit from the labor without having to pay developers. However, mixed strategies are best for reinforcement learning technologies. Indeed, patents, trade secrets, copyrights, and open source data are all vital features for a robust property portfolio.

Appendices

Appendix A. Notation Summary

Notation	Meaning
$Pr(S)$	Probability of sequence
$Pr(T)$	Language model of the target language
$Pr(T S)$	Translation model
V^π	Value function
γ	Discount factor
R	Reward
Q	Q value
π	Policy
a	Action
s	State

Appendix B. Global Statistics

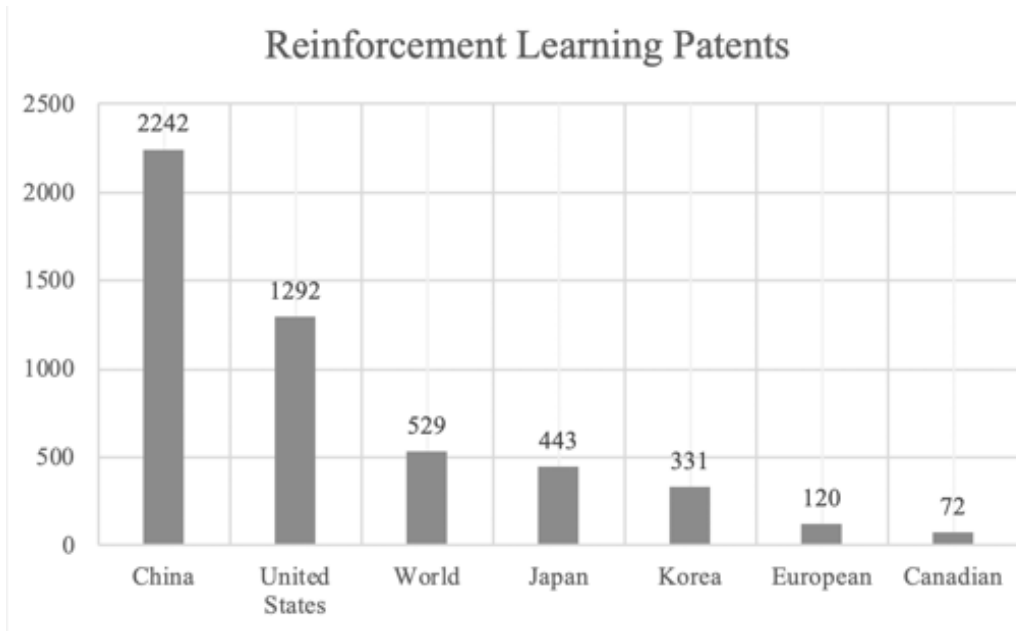


Figure 9²⁵¹

²⁵¹ As of July 30, 2020.

Appendix C. Top Ten Most Valuable Patents

Rank ²⁵²	Number	Title	Value (USD) ²⁵³	Holder
1	6,572,542	System and method for monitoring and controlling the glycemic state of a patient	\$33,617,905.81	Medtronic, Inc.
2	10,467,274	Deep reinforcement learning-based captioning with embedding reward	\$33,516,953.65	Snap Inc.
3	9,104,905	Automatic analysis of individual preferences for attractiveness	\$32,171,404.34	Emotient, Inc.
4	8,231,542	System for analyzing thermal data based on breast surface temperature to determine suspect conditions	\$31,941,887.56	Lifeline Biotechnologies, Inc.
5	10,249,206	Systems and methods for providing information incorporating reinforcement-based learning and feedback	\$31,613,171.74	D2L Corporation
6	7,839,292	Real-time driving danger level prediction	\$31,507,535.43	NEC Laboratories America, Inc.
7	10,146,286	Dynamically updating a power management policy of a processor	\$30,787,709.36	Intel Corporation
8	9,679,258	Methods and apparatus for reinforcement learning	\$30,571,846.46	Google Inc.
9	9,741,336	System and method for generating manually designed and automatically optimized spoken dialog systems	\$30,505,181.84	Nuance Communications, Inc.
10	9,047,225	Dynamic selection of data replacement protocol for cache	\$30,234,053.62	EMC Corporation
HM	10,032,281	Multi-scale deep reinforcement machine learning for N-dimensional segmentation in medical imaging	\$29,345,457.34	Siemens Healthcare GmbH
HM	10,503,174	Method and device for optimized resource allocation in autonomous driving on the basis of reinforcement learning using data from lidar, radar, and camera sensor	\$29,169,576.93	StradVision, Inc.

²⁵² Ranking for most valuable U.S. Patents for reinforcement learning technology.

²⁵³ All values were calculated using artificial intelligence algorithms.

Appendix D. Transatlantic Translations

Number	EP3584649A1
English	Facility state determination device, facility state determination method, and facility management system
Français	Dispositif de détermination de l'état d'une installation, procédé de détermination d'état d'une installation et système de gestion d'installations
Deutsche	Einrichtungszustandsbestimmungsvorrichtung, einrichtungszustandsbestimmungsverfahren und einrichtungsverwaltungssystem
Number	EP3579154A1
English	Reinforcement learning for user behavior
Français	Apprentissage par renforcement pour comportement d'utilisateur
Deutsche	Verstärkungslernen für das Nutzerverhalten
Number	EP3579189A1
English	Adaptive nonlinear optimization of shape parameters for object localization in 3D medical images
Français	Optimisation linéaire adaptative de paramètres de forme pour la localisation d'objets dans des images médicales en 3D
Deutsche	Adaptive nichtlineare Optimierung von Formparametern zur Objektortung in medizinischen 3D-Bildern
Number	EP3575172A1
English	Adaptive longitudinal control using reinforcement learning
Français	Commande longitudinale adaptative utilisant l'apprentissage par renforcement
Deutsche	Adaptive Längsregelung durch Verstärkungslernen
Number	EP3576038A1
English	Trade platform with reinforcement learning
Français	Plateforme d'échange avec apprentissage par renforcement
Deutsche	Handelsplattform mit Verstärkungslernen
Number	EP3573371A1
English	Optimizing a wireless network that comprises range extenders
Français	Optimisation d'un réseau sans fil comprenant des prolongateurs de portée
Deutsche	Optimierung eines drahtlosen Netzwerks mit Entfernungserweiterungen
Number	EP3576023A1
English	Trade platform with reinforcement learning network and matching engine
Français	Plateforme de commerce comportant un réseau d'apprentissage par renforcement et un moteur d'appariement
Deutsche	Handelsplattform mit Verstärkungslernnetzwerk und Abgleichungsmotor
Number	EP3567539A1
English	Method and system for orchestrating multi-party services using semi-cooperative nash equilibrium based on artificial intelligence, neural network models, reinforcement learning and finite-state automata
Français	Procédé et système d'orchestration de services multipartites utilisant l'équilibre de Nash semi-coopératif basé sur l'intelligence artificielle, modèles de réseau neuronal, Apprentissage par renforcement et automates à l'état fini
Deutsche	Verfahren und System zur Orchestrierung von Mehrparteiendiensten unter Verwendung eines semi-kooperativen Nashgleichgewichts basierend auf künstlicher Intelligenz, neuronale Netzwerkmodelle, Verstärkungslernen und Zustandsautomaten
Number	EP3557489A1

English	Energy optimization in operation of rail vehicle
Français	Optimisation dénergie lors du fonctionnem ent d'un véhicule ferroviaire
Deutsche	Energieoptimierung im betrieb eins zugfahrzeugs
Number	EP3553710A1
English	Artificial intelligence optimized telecommunications systems
Français	Systèmes de telecommunication optimisés par intelligence artificielle
Deutsche	Mit künstlicher intelligenz optimierte telekommunikationssysteme
Number	EP3564861A1
English	Vision-based sample-efficient reinforcement learning framework for autonomous driving
Français	Structure d'apprentissage par renforcement efficace en matière d'échantillons basée sur la vision pour conduit autonome
Deutsche	Sichtbasierter probeneffizienter verstärkungslernen-rahmen für autonomes fahren
Number	EP3543918A1
English	Reinforcement learning method
Français	Procédé d'apprentissage par renforcement
Deutsche	Verstärkungslernverfahren
Number	EP3530463A1
English	Apparatus and method of generating control parameter of screen printer
Français	Appareil et procédé de génération d'un paramètre de commande d'une machine de sérigraphie
Deutsche	Vorrichtung und verfahren zur erzeugung eines steuerparameters einer siedruckmaschine
Number	EP3528463A1
English	An artificial intelligence cyber security analyst
Français	Analyste de cybersécurité en intelligence artificielle
Deutsche	Cybersicherheitsanalysator mit künstlicher intelligenz
Number	EP3525136A1
English	Distributed machine learning system
Français	Système distribué d'apprentissage automatique
Deutsche	System zum verteilten maschinellen lernen
Number	EP3515038A1
English	Autonomous reconfigurable virtual sensing system for cyber-attack neutralization
Français	Système de detection virtuelle reconfigurable autonome pour la neutralisation des cyber-attaques
Deutsche	Autonomes rekonfigurierbares virtuelles sensorsystem zur neutralisierung von cyberangriffen
Number	EP3501385A1
English	System and method for determining a subject's stress condition
Français	Système et procédé permettant de déterminer l'état de stress d'un sujet
Deutsche	System und verfahren zur bestimmung des stresszustands einer person
Number	EP3483895A1
English	Detecting and classifying medical images based on continuously-learning whole body landmarks detections
Français	Détection et classification d'images médicales sur la base de détections de repère de corps entier par Apprentissage continu
Deutsche	Detektion und klassifizierung von medizinischen bildern auf basis von kontinuierlich lernenden ganzkörperlandmarkendetektionen
Number	EP3483785A1

English	System and method for guiding social interactions
Français	Système et Procédé pour guider des interactions sociales
Deutsche	System und verfahren zum leiten sozialer interaktionen
Number	EP3467718A1
English	Machine learning system
Français	Système d'apprentissage par machine
Deutsche	System zum maschinellen lernen
Number	EP3467662A1
English	Verification of applications that utilize artificial intelligence
Français	Vérification d'applications d'intelligence artificielle
Deutsche	Prüfung der künstliche intelligenz anwendungen
Number	EP3459810A1
English	Method for predicting failure of a sensor
Français	Procédé permettant de prédire une défaillance de capteur
Deutsche	Verfahren zur vorhersage des ausfalls eines sensors
Number	EP3454158A1
English	Method and device for navigation of an industrial truck
Français	Procédé et dispositif de navigation d'un chariot de manutention
Deutsche	Verfahren und vorrichtung zur navigation eines flurförderzeugs
Number	EP3446820A1
English	Wire electrical discharge machining method
Français	Dispositif d'usinage à décharge électrique de câble
Deutsche	Drahterosionsbearbeitungsverfahren
Number	EP3444824A1
English	Detecting and classifying medical images based on continuously-learning whole body landmarks detections
Français	Détection et classification d'images médicales sur la base de détections de repère de corps entier par Apprentissage continu
Deutsche	Detektion und klassifikation von medizinischen bildern auf basis von kontinuierlich lernenden ganzkörperlandmarkendetektionen
Number	EP3657404A1
English	Machine learning and object searching method and device
Français	Procédé et dispositif d'apprentissage automatique et de recherche d'object
Deutsche	Verfahren und vorrichtung für maschinellen lernen und objektsuche
Number	EP3428925A1
English	Method and system for clinical decision support with local and remote analytics
Français	Procédé et système de support de décision clinique comportant des analyses locales et distantes
Deutsche	Verfahren und system für klinische entscheidungsunterstützung mit lokaler und entfernter analytik
Number	EP3422130A1
English	Method and system for autonomously operating an aircraft
Français	Procédé et système de fonctionnement autonome d'un aéronef
Deutsche	Verfahren und system zur autonomen steuerung eines flugzeugs
Number	EP3407194A2
English	Method for the deployment of distributed fog computing and storage architectures in robotic modular components
Français	Procédé de déploiement d'architectures de calcul et de stockage de broillard réparties dans des composants modulaires robotiques

Deutsche	Verfahren zum einsetzen von verteiltem fog-computing and speicherarchitekturen in robotische modulare komponenten
Number	EP3407264A1
English	Listen, interact, and talk: learning to speak via interaction
Français	Écouter, interagir et discuter: apprendre à parler via une interaction
Deutsche	Zuhören, interagieren und sprechen: sprechen lernen durch interaktion
Number	EP3396598A2
English	Method and user interface for managing and controlling power in modular robots and apparatus therefor
Français	Procédé et interface utilisateur pour la gestion et la régulation de la puissance dans des robots modulaires et appareil correspondant
Deutsche	Verfahren und benutzerschnittstelle zur verwaltung und steuerung von energie in modularen robotern und vorrichtung dafür
Number	EP3385884A1
English	Method for recognizing an object of a mobile unit
Français	Procédé de reconnaissance d'un objet d'une unite module
Deutsche	Verfahren zur erkennung eines objekts einer mobilen einheit
Number	EP3584750A1
English	Image processing device, information processing method, and program
Français	Dispositif de traitement d'informations, Procédé de traitement d'information et programme
Deutsche	Informationsverarbeitungsvorrichtung, informationsverarbeitungsverfahren und programm
Number	EP3579536A1
English	Image processing system, optical sensor, and learning device
Français	Système de traitement d'image, capteur optique et dispositif d'apprentissage
Deutsche	Bildverarbeitungssystem, optischer sensor und lernvorrichtung
Number	EP3399501A1
English	Multi-scale deep reinforcement machine learning for n-dimensional segmentation in medical imaging
Français	Apprentissage machine de reinforcement profonde multi-échelles pour segmentation n-dimensionnel en imagerie médicale
Deutsche	Multiskalares tiefenverstärkendes maschinenlernen für n-dimensionale segmentierung in der medizinischen bildgebung
Number	EP3561740A1
English	Learning device and learning method
Français	Dispositif d'apprentissage et procédé d'apprentissage
Deutsche	Lernvorrichtung und lernverfahren
Number	EP3339880A1
English	Adaptive radar system
Français	Système radar décentralisé
Deutsche	Adaptives radarsystem
Number	EP3340118A1
English	Traced-based neuromorphic architecture for advanced learning
Français	Architecture neuromorphique à base de trace pour l'apprentissage avancé
Deutsche	Ablaufverfolgungsbasierte neuromorphe architektur für fortgeschrittenes lernen
Number	EP3340119A1
English	Rapid competitive learning techniques for neural networks
Français	Techniques d'apprentissage concurrent el rapide pour réseaux neuronaux

Deutsche	Schnelle kompetitive lerntechniken für neuronale netze
Number	EP3553711A1
English	Information processing device and method, and program
Français	Procédé et dispositif de traitement d'informations, et programme
Deutsche	Informationsverarbeitungsvor richtung und-verfahren und programm
Number	EP3319016A1
English	Control systems using deep reinforcement learning
Français	Systèmes de commande utilisant un apprentissage par renforcement
Deutsche	Steuerungssysteme mit tiefenverstärkungslernen
Number	EP3531369A1
English	Information processing device and information processing method
Français	Dispositif de traitement de l'information et procédé de traitement de l'information
Deutsche	Informationsverarbeitungsvor richtung und informationsverarbeitungsverfahren
Number	EP3279820A1
English	Medical scanner teaches itself to optimize clinical protocols and image acquisition
Français	Scanner médical á autoenseignement afin d'otimiser les protocoles cliniques et l'acquisition d'images
Deutsche	Medizinischer scanner lehrt sich selbst die optimierung klinischer protokollow und bilderfassung
Number	EP3295611A1
English	Early warning reconommendation system for the proactive management of wireless broadband networks
Français	Système d'alerte précoce et de recommandation pour la gesti3n proactive de réseaux sans fil á haut débit
Deutsche	Frühwarn-und empfehlungssystem zur proaktiven verqaltung von drahtlosbreitbandnetzwerken
Number	EP3270095A1
English	System and method for surface inspection
Français	Système et procédé d'inspection de surface
Deutsche	System und verfahren zur oberflächenprüfung
Number	EP3246875A2
English	Method and system for image registration using an intelligent artificial agent
Français	Procédé et système d'enregistrement d'image á l'aide d'un agent artificiel intelligent
Deutsche	Verfahren und system zur bildregistrierung mit einem intelligenten künstlichen mittel
Number	EP3452984A1
English	Medical Atlas Registration
Français	Recalage d'atlas médical
Deutsche	Medizinische atlasregistrierung
Number	EP3435296A1
English	Information processing device
Français	Procédé de commande d'un système de refroidissement avec pussance de refroidissement variable, et le système de refroidissement
Deutsche	Informationsverarbeitungsvorrichtung
Number	EP3208681A1
English	Control method for a cooling system with variable cooling power and cooling system
Français	Procédé de commande d'un système de refroidissement avecpuissance de refroidissement variable, et le systeme de refroidissement
Deutsche	Steuerungsverfahren für ein kühlssystem mit variabler kühlleistung sowie kühlssystem
Number	EP3142033A1

English	Physiology-driven decision support therapy planning
Français	Support de décision guidé par la physiologie pour une planification de thérapie
Deutsche	Physiologiegesteuerte entscheidungsunterstützung zur therapieplanung
Number	EP3136304A1
English	Methods and systems for performing reinforcement learning in hierarchical and temporally extended environments
Français	Procédés pour effectuer l'apprentissage par renforcement dans des environnements hiérarchiques et prolongés dans la durée
Deutsche	Vergahren und systeme zur durchführung von verstärkungslernen in hierarchischen und zeitlich erweiterten umgebungen
Number	EP3079106A2
English	Selecting reinforcement learning actions using goals and observations
Français	Actions d'apprentissage par renforcement de selection utilisant des objectifs et des observations
Deutsche	Auswahl von bestärkenden lernmassnahmen unter verwendung von zielen und beobachtungen
Number	EP3075496A1
English	Method for improving operation of a robot
Français	Procédé permettant d'améliorer le fonctionnement d'un robot
Deutsche	Verfahren zur verbesserung des betriebs eines roboters
Number	EP3172434A1
English	System and methodology for detecting structural damages of wind turbine rotor blades
Français	Système et méthodologie permettant de détecter des dégâts structurels de pales de rotor d'éolienne
Deutsche	System und Verfahren zur erkennung struktureller schäden von windturbinenrotoblättern
Number	EP3129839A1
English	Controlling a target system
Français	Commande d'un systeme cible
Deutsche	Steuerung eines zielsystems
Number	EP2947035A1
English	Method for determining the load on a working machine and working machine, in particular a crane
Français	Procédé de determination de la charge d'une machine de travail et machine de travail, en particulier une grue
Deutsche	Verfahren zur bestimmung der aufgenommenen last einer arbeitsmaschine sowie arbeitsmaschine, insbesondere kran
Number	EP3117274A1
English	Method, controller, and computer program product for controlling a target system by separately training a first and a second recurrent neural network models, which are initially trained using operational data of source systems
Français	Procédé, dispositif de commande et produit-programme d'ordinateur permettant de commander un système cible par l'entraînement distinct de premier et second modèles de réseaux neuronaux récurrents qui sont initialement entraînés au moyen de données opérationnelles de systèmes sources
Deutsche	Verfahren, steuergerät und computerprogrammprodukt zur steuerung eines zielsystems durch separates trainieren eines ersten und eines zweiten wiederkehrenden models eines neuralen netzes durch anfängliches trainieren mit betriebsdaten von quellsystemen

Number	EP3055813A1
English	Methods and apparatus for reinforcement learning
Français	Procédés et appareil d'apprentissage par renforcement systèmes
Deutsche	Verfahren und vorrichtung für bestärkendes lernen
Number	EP3022610A2
English	Method and device for controlling the behavior of systems
Français	Procédé et dispositif de commande du comportement de
Deutsche	Verfahren und vorrichtung zur verhalten steuerung von systemen
Number	EP2784710A2
English	Method and system for validating personalized account identifiers using biometric authentication and self-learning algorithms
Français	Procédé et système pour des identifiants de compte personnalisés au moyen d'une authentification biométrique et d'algorithmes d'auto-apprentissage
Deutsche	Verfahren und system zur validierung persönlicher konto-identifier unter verwendung der biometrischen authentizierung und lernfähige algorithmen
Number	EP2747357A1
English	Robust content-based solution for dynamically optimizing multi-user wireless multimedia transmission
Français	Solution robuste à base de contenu pour optimiser dynamiquement une transmission multimédia sans fil multi-utilisateurs
Deutsche	Robuste inhaltsbasierte lösung für die dynamisch optimierende Multiuser-Funk-Multimedien-Übertragung
Number	EP2790329A1
English	Card-type wireless transceiver for a vehicle, and method for manufacturing same
Français	Appareil émetteur-récepteur sans fil du type à carte, pour un véhicule, et procédé pour sa fabrication
Deutsche	Kartenförmiger drahtloser sende-empfänger für ein fahrzeug und verfahren zu seiner herstellung
Number	EP2801395A1
English	Golf club for teaching or learning golf
Français	Golfschläger zum lehren oder erlernen von golf
Deutsche	Canne de golf pour l'enseignement ou l'apprentissage du golf
Number	EP2697695A2
English	Method for the computer-supported generation of a data-driven model of a technical system, in particular of a gas turbine or wind turbine
Français	Procédé pour générer, de manière assistée par ordinateur, un modèle piloté par des données d'un système technique, notamment d'une turbine à gaz ou d'une éolienne
Deutsche	Verfahren zur rechengestützten generierung eines datengetriebenen modells eines technischen systems, insbesondere einer gasturbine oder windturbine
Number	EP2501118A2
English	Method and apparatus for identifying a conference call from an event
Français	Procédé et appareil pour identifier un appel en conférence à partir d'un enregistrement d'événements
Deutsche	Verfahren und vorrichtung zur identifizierung eines konferenzerufes aus einer ereignisaufzeichnung
Number	EP2641214A1
English	Electronic synapses for reinforcement learning
Français	Synapses électronique es à l'apprentissage renforcé
Deutsche	Elektronische synapsen zum verstärkungslernen

Number	EP2609792A2
English	Automatic configuration of a lighting system
Français	Configuration automatique d'un system d'éclairage
Deutsche	Automatische configuration eines beleuchtungssystems
Number	EP2585248A1
English	Method for controlling a laser processing operation by means of a reinforcement learning agent and laser material processing head using the same
Français	Procédé de commande d'une opération de traitement laser au moyen d'un agent d'apprentissage par renforcement et tête de traitement de matériau au laser utilisant ce procédé
Deutsche	Verfahren zur steuerung eines laserbearbeitungsvorgangs anhand eines verstärkenden lernagenten und lasermaterialbearbeitungskopf damit
Number	EP2386987A1
English	A method of reinforcement learning, corresponding computer program product, and data storage device therefor
Français	Procédé de renforcement de l'apprentissage, produit de programme informatique correspondant et dispositif de stockage de données correspondant
Deutsche	Verfahren zum bestärkenden Lernen, entsprechendes Computerprogrammprodukt und Datenspeichervorrichtung dafür
Number	EP2381393A1
English	A method of reinforcement learning, corresponding computer program product, and data storage device therefor
Français	Procédé de renforcement de l'apprentissage, produit de programme informatique correspondant et dispositif de stockage de données correspondant
Deutsche	Verfahren zum bestärkenden lernen, entsprechendes computerprogrammprodukt und datenspeichervorrichtung dafür
Number	EP2381394A1
English	A method of reinforcement learning, corresponding computer program product, and data storage device therefor
Français	Procédé de renforcement de l'apprentissage, produit de programme informatique correspondant et dispositif de stockage de données correspondant
Deutsche	Verfahren zum bestärkenden Lernen, entsprechendes Computerprogrammprodukt und datenspeichervorrichtung dafür
Number	EP2519861A1
English	Method for the computer aided control of a technical system
Français	Verfahren zur rechnergestützten steuerung und/oder regelung eines technischen systems
Deutsche	Procédé de commande et/ou de régulation assistée par ordinateur d'un système technique
Number	EP2456592A1
English	Laser machining head and method of compensating for the change in focal position of a laser machining head
Français	Tête d'usinage au laser et procédé permettant de compenser la variation de position du foyer pour une tête d'usinage au laser
Deutsche	Laserbearbeitungskopf und Verfahren zur compensation der fokuslagenänderung beieinem laserbearbeitungskopf
Number	EP2379786A1
English	Method for controlling a laundry distribution mode of a domestic appliance for caring for laundry items

Français	Procédé pour commander une fonction de repartition du linge dans un appareil ménager destiné à l'entretien du linge
Deutsche	Verfahren zum stuern eines wäscheverteils eines haushaltgeräts zur pflege von wäschestücken
Number	EP2206023A1
English	Method for computer-assisted exploration of states of a technical system
Français	Procédé pour l'exploration assistée par ordinateur des états d'un système technique
Deutsche	Verfahren zur rechnergestützten exploration von zuständen
Number	EP2108139A1
English	Method for the computer aided control and/or regulation of a technical system, particularly a gas turbine
Français	Procédé de regulation et/ou de contrôle assisté par ordinateur d'un système technique, en particulier d'une turbine à gaz
Deutsche	Verfahren zur rechnergestützten regelung und/oder steuerung eines technischen systems, insbesondere einer gasturbine
Number	EP2097793A1
English	Method for the computer-assisted control and/or regulation of a technical system
Français	Procédé de commande et/ou de régulation d'un système technique assistés par ordinateur
Deutsche	Verfahren zur rechnergestützten steuerung und/oder regelung eines technischen systems
Number	EP1764933A1
English	Aeronautical communication system
Français	Système de communications aéronautiques
Deutsche	Aeronautisches kommunikationssystem
Number	EP2360629A2
English	Device for the autonomous bootstrapping of useful information
Français	Dispositif d'amorçage autonome d'informations utiles
Deutsche	Vorrichtung zum automaitischen urladen von nützlichen informationen
Number	EP1650672A2
English	A neural network element with reinforcement/attenuation learning
Français	Élément de réseau neuronal avec apprentissage de renforcement/atténuation
Deutsche	Element eines neuronalen Netzes mit bestärkendem/abschwächendem Lernen
Number	EP1677668A1
English	Apparatus and method for determining physical location
Français	Appareil et procede pour determiner un et at physiologique
Deutsche	Gerät und Verfahren zur bestimmung eines physiologischen zustands
Number	EP1528464A2
English	Proactive user interface including evolving agent
Français	Interface utilisateur proactive ayant un agent évolutif
Deutsche	Proaktive Benutzerschnittstelle mit entstehendem agenten
Number	EP1416744A1
English	Multicriterial dynamic routing in communication networks
Français	Routage dynamique selon des critères multiples dans des réseaux de communication
Deutsche	Dynamische Leitweglenkung mit mehreren Kriterien in Kommunikationsnetzwerken
Number	EP2325294A2
English	Visual-servoing optimal microscopy
Français	Microscopie optique à asservissement visuel
Deutsche	Adaptives learnsystem und-verfahren

Number	EP1287488A1
English	Adaptive learning system and method
Français	Systeme et technique d'apprentissage adaptatif
Deutsche	Adaptives learnsystem und-verfahren
Number	EP1016981A1
English	Agent learning machine
Français	Machine didactique pour agents
Deutsche	Didaktische maschine für agenten
Number	EP0935202A1
English	Hardware or software architecture implementing self-biased conditioning
Français	Architecture matérielle ou de logiciel avec conditionnement autopolarisé
Deutsche	Hardware-oder Software-Architektur mit selbstvorgespannter Konditionierung
Number	EP0885420A1
English	Method of training a neural network
Français	Procede d'apprentissage pour unreseau neuronal
Deutsche	Verfahren zum trainieren eines neuronalen netzes
Number	EP0462916A2
English	Neural network model for reaching a goal state
Français	Modèle à réseau neuronal pour atteindre un état désiré
Deutsche	Neuronalnetzmodell
Number	EP0434423A2
English	A system for learning an external evaluation standard
Français	Système pour apprendre un standard d'évaluation externe
Deutsche	System zum lernen eines externen auswertungsstandards
Number	EP0157080A2
English	Probabilistic learning element
Français	Elément autoadaptatif à fonction de probabilié
Deutsche	Selbstlernendes element mit wahrescheinlichkeitsfunktion
Number	EP0159463A2
English	Probabilistic learning system
Français	Système autoadaptif à fonction de probabilité
Deutsche	Selbstlernendes system mit wahrscheinlichkeitsfunktion