This work has been supported by Autodesk Research and The Studio for Creative Inquiry, and was made possible in part with additional sponsorship from Autodesk, the Autodesk Artist in Residence program, the Design Museum, and ABB Robotics.

Human-Centered Interfaces for Autonomous Fabrication Machines

MADELINE GANNON

CMU-SOA-18-01
January 2018

COMPUTATIONAL DESIGN

School of Architecture
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15217

mgannon@andrew.cmu.edu
https://atonaton.com

COMMITTEE

Ramesh Krishnamurti (Chair), Carnegie Mellon University
Golan Levin,Carnegie Mellon University
Tovi Grossman, Autodesk Research
Aisling Kelliher, Virginia Tech

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Keywords

Interaction, design, human-centered, interface, gesture, 3D modeling, 3D printing, on-body interface, on-body design, on-body fabrication, wearables, human-computer interaction, human-robot interaction, industrial robot.

Abstract

Computer-controlled fabrication machines have been an essential part of industrial infrastructure since the early 1960s. The unique abilities of these machines – their speed, precision, strength, endurance, and programmability – have long provided a strategic advantage for factory automation. However, today we are witnessing the transition from *automated* to *autonomous* systems of production. Now, instead of being restricted to short, repetitive, pre-programmed tasks, fabrication machines are gaining the ability to dynamically see and respond to their changing environment. While this transition represents a significant advancement for manufacturing, it also presents a newfound opportunity to explore human-centered interaction design with these large, potentially dangerous, non-humanoid machines.

In this dissertation, I demonstrate the potential for rebalancing systems of automation to be more inclusive of people. My research examines how to combine intelligent sensing with well-designed interfaces, so that the underlying control framework of a fabrication machine can better understand a human-counterpart. I present three interactive systems that progressively embody a fabrication machine with this contextual information: with *Reverb*, I develop a framework for embedding machine knowledge into interactive, semi-autonomous geometry; in *Tactum*, I demonstrate how to adapt this intelligent, fabrication-aware geometry to dynamically changing physical environments; with *ExoSkin*, I examine the technical challenges of direct, close-quarter interaction with fabrication machines.

Finally, my work culminates with *Mimus*, a 1,200kg industrial robot that I transformed into a living, breathing mechanical creature. *Mimus* synthesizes innovations and techniques first developed in *Reverb*, *Tactum*, and *ExoSkin* to illustrate new interaction possibilities when coexisting with autonomous, attentive machines. It re-examines the unique affordances of an industrial robot, and illustrates how an existing tool of automation can be reconfigured to have more meaningful interactions with people. This body of work demonstrates the potential for human-centered interfaces to combine the unique abilities of people and machines in ways to transcend one another's limitations. In doing so, my research aspires to show how our systems of automation can be reconfigured to enhance, augment, and expand human capabilities – not replace them.

Acknowledgements

I am forever grateful to the members of my PhD committee, whose guidance, mentorship, and friendship have been fundamental to the development of my research practice. I am fortunate to have a committee made up of world renown experts across diverse disciplines, including Computational Geometry, Design, Human-Computer Interaction, and New Media Art. Their thoughtful supervision (and strong personalities) have helped me identify new, interdisciplinary territories to explore.

Professor Ramesh Krishnamurti welcomed me as a bright-eyed, young Masters student into his Computation Design program at Carnegie Mellon University back in 2010. When he invited me to join the PhD program in 2012, I had no idea where my research path would lead me. However, I was bolstered by his confidence and trust that I would eventually find my footing. Thank you for seeing something in me that I did not, and for patiently tolerating my creative interpretation of the rules.

I am grateful to *Tovi Grossman* for taking a chance on me, and hiring someone from a college of fine arts as a User Interface Researcher. Thank you for welcoming me into the HCI community, and for your tireless support throughout the years. I wish you best for your new role in academia: I am thrilled for the next generation of young researchers who will be working with you, and excited for the new paths you'll pave with them.

Thank you to *Aisling Kelliher*, who first showed me how rigorous academic research can be enhanced with thoughtful artistic inquiry. You have been an incredible role model throughout my time as a PhD candidate, and I am grateful to you for tempering my wild ambitions towards more focused and productive ends.

Golan Levin has been my loudest champion while at Carnegie Mellon University. He is a singular force that has pushed my work to reach further and be less boring. Golan is most responsible for my ability to clearly and thoughtfully communicate my work. I am honored to be your first doctorate student, and the first PhD to come out of the STUDIO for Creative Inquiry. I hope my work here will open more opportunities for future students to benefit from your guidance.

Thank you to my unofficial mentors and spirit guides – especially *Mark Gross*, *Gill Wildman*, and *Nick Durrant* – for patiently incubating my raw enthusiasm as a young designer. Your generosity has helped me mature into a strategic thinker, designer, and researcher capable of impacting the world. I am also indebted to my coauthors – *George Fitzmaurice* and *Tovi Grossman* – my collaborators – *Zack Jacobson-Weaver* and *Dan Moore* – and my colleagues, whose considerable efforts have amplified the achievements of this body of work. This holds especially true for my development team with *Mimus* – *Julián Sandoval*, *Kevyn McPhail*, and *Ben Snell* – whose trust, support, and sheer power-of-will helped make this ambitious project possible.

I am fortunate to have had many sponsors who have empowered me to dream, develop, and demonstrate my visions of the future to the world. Thank you to Autodesk, ABB Robotics, the Design Museum, and the Frank-Ratchye STUDIO for Creative Inquiry for amplifying my voice.

Autodesk has been a steadfast supporter over the years, and I feel lucky to have met and worked with so many talented, interesting people across their organization. Thank you to my friends at Autodesk Research in Toronto, the Autodesk Artist in Residence Program and Applied Research Lab in San Francisco, and the Autodesk Builder in Residence Program in Boston for always pushing me to create more future-forward work.

The Frank-Ratchye STUDIO for Creative Inquiry has been my home away from home during my tenure at Carnegie Mellon University. Thank you to Tom Hughes, Linda Hager, and Carol Hernandez whose patience, support, and friendship I will always cherish. To Sarah Ratchye and Ed Frank: I can't overstate the impact that your support has had on my burgeoning career, and those of countless others. Thank you for empowering the STUDIO to explore new modes of arts research, and for fostering a welcoming, supportive home to a herd of black sheep.

Finally, this body of work would not have been possible without my family. I dedicate this dissertation to my mother, who gave me the curiosity and confidence to reimagine the world, to my father, who gave me the work ethic and ambition to always climb the next tallest mountain, and to my husband, who gave me the unconditional love I needed to leap into the unknown. *I love you all.*

Table of Contents

Chapter 1

Introduction

Living side-by-side with robots may seem like a vision of a science fiction future, but it has been a reality for people working in manufacturing for over half a century. Computer-controlled fabrication machines have been an essential part of automation infrastructure since the early 1960s. The unique abilities of these machines – their speed, precision, strength, endurance, and programmability – have long been harnessed to produce commodities faster, cheaper, and safer than most non-automated processes. Moreover, in driving down the cost of production, fabrication machines acted as a strategic advantage towards positioning the United States as a global industrial super-power in the mid-twentieth century [Noble 1984].

Today, we are witnessing the transition from *automated* to *autonomous* systems of production. Increased access to advanced sensing, machine learning (ML), and artificial intelligence (AI) techniques are augmenting existing machines with new levels of understanding for their surroundings. Now, instead of being restricted to short, repetitive, pre-programmed tasks, fabrication machines are gaining the ability to dynamically see and respond to their changing environment without the need for human involvement.

While a transition towards autonomous machines is a significant advancement for manufacturing, there are concerns over the human cost of this technological progress. For example, up to 5 million jobs are projected to be displaced by 2020 due the rate of global adoption for robotic automation [World Economic Forum 2016]. Moreover, researchers and policy makers are cautioning against the potential disruptive socioeconomic impact brought on by the rapidly accelerating intelligence and independence of autonomous machines [Knight 2014; Pew Research 2014].

Despite such well-founded anxieties, these recent advancements present a newfound opportunity to explore more human-centered interactions with autonomous fabrication machines. Usability and safety constraints have long prevented robust exploration of these machines for

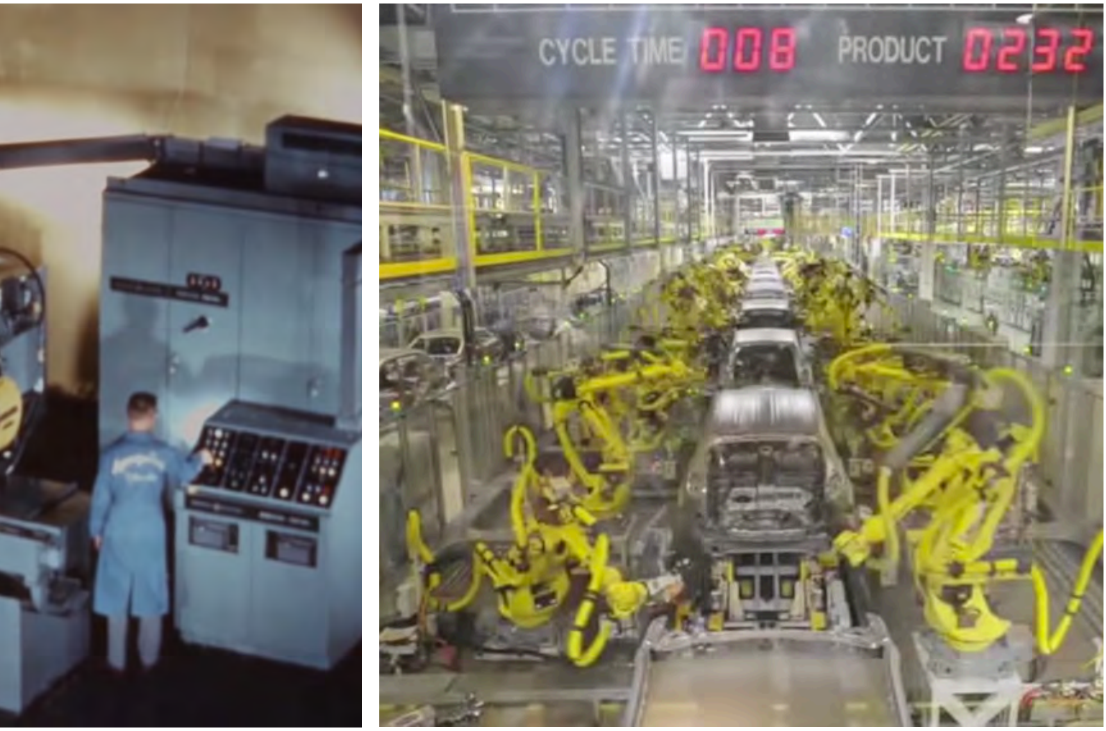

FIGURE 1

People and machines have long co-existed in manufacturing settings. However, this dynamic is changing as these machines increasingly gain intelligence and autonomy.

applications outside of the context of manufacturing. A lack of intelligent sensing has made it quite dangerous for these machines to operate near irregular or unpredictable objects, such as people. However, in this dissertation I show how thoughtful interaction design, when combined with intelligent sensing, can reconfigure the underlying control systems of fabrication machines to be more inclusive of people. This body of work demonstrates the potential for well-designed interfaces to combine the unique abilities of people and machines in ways to transcend one another's limitations. In doing so, this research seeks to inspire an optimistic vision for rebalancing our systems of automation to enhance, augment, and expand human capabilities – not replace them.

Organization

For the rest of this chapter, I provide the necessary context to understand the scope, motivations, and outlook for my inquiry into human-centered interfaces for autonomous fabrication systems. I begin by introducing the socio-economic impact that fabrication machines have had in shaping modern society. I then discuss the current and future challenges that these machines now pose. Lastly, I highlight strategies for mitigating these negative effects of automation.

Chapter 2 provides a summary of related works that have been most influential to my research. My work pulls from many disciplinary topics – including design, new media art, human-computer interaction, and human-robot interaction. However, in this section, I focus on the most relevant systems in the fields of gestural user interfaces, CAD/CAM, computer vision, human-centered design, 3D printing, and industrial robotics.

The body of my research – Chapters 3, 4, and 5 – details my process of discovering the conceptual, technical, and interaction possibilities of engaging with intelligent and independent machines. Each chapter documents a *progressive embodiment* of contextual information into the underlying control framework of a fabrication machine: through *Reverb*, I develop a framework for embedding machine knowledge into interactive, semi-autonomous geometry; through *Tactum*, I demonstrate how to adapt this intelligent, fabrication-aware geometry to a dynamically changing physical environment; with *ExoSkin*, I examine the technical challenges of working in close-quarters directly with a fabrication machine.

Chapter 3 presents *Reverb*: a gestural 3D modeling system that integrates autonomous behaviors into its *virtual geometry*. This chapter focuses most heavily on examining the unique affordances of intelligent *geometry*. It highlights techniques for embedding the technical knowledge of a fabrication process into animate digital forms, and it discusses how offloading this cognitive overhead can make such systems easier and more intuitive to use for designers.

Chapter 4 presents *Tactum*: an on-body 3D modeling system for designing wearables directly on the skin. This chapter expands on the insights learned through *Reverb* to focus on autonomous behaviors in *contextual interfaces*. It illustrates how intelligent, fabrication-aware geometry can come out of the computer and occupy a physical space. Finally, it discusses the benefits of digital design and fabrication systems that can adapt to dynamically changing environments.

Chapter 5 presents *ExoSkin*: an intelligent guidance system for fabricating wearables directly on the body. It continues this thread to examine how a designer and a fabrication machine can work in tangible, close cooperation with one another. This chapter details how machines built around human affordances can expand digital design and fabrication systems into previously analog domains, and it discusses how machines integrated with context-aware geometry and interfaces can make the line between *tool* and *collaborator* more ambiguous.

Chapter 6 synthesizes the techniques developed and refined through *Reverb*, *Tactum*, and *ExoSkin* to examine ways of coexisting with intelligent, attentive, animate machines. This chapter presents *Mimus*, a 1,200kg industrial robot transformed into a living, breathing mechanical creature. It demonstrates how a well-designed, human-centered interface can allow many non-technical users to successfully engage with a large, complex fabrication machine in a highly public setting. The intimate connection between Mimus and her visitors illustrates aspirational new relationships that may form between people and intelligent fabrication machines.

Finally, Chapter 7 summarizes my key findings and contributions, outlines optimistic paths towards future work, and offers some closing thoughts.

Scope & Methodology

This research is primarily focused on the *interaction challenges* of engaging with autonomous fabrication machines. The interactive systems I present often showcase physical artifacts that help illustrate how human-centered interfaces can foment better ways of communicating with machines that make things. However, in this research, *what* the fabrication machine makes is less important than the interaction between the designer, the environment, and the machine.

Furthermore, in this dissertation fabrication machines are considered proxies for other large, potentially dangerous, non-humanoid robots. These ubiquitous, versatile, and reliable machines are a mature robotic platform for exploring many of the safety and usability issues of newer robots that are only now beginning to join us in the built environment. Moerover, as roboticists continue to invent novel machines may join us in our homes, offices, cities, or skies, our long-standing rapport with fabrication machines can offer a valuable perspective into how human-robot relations can shift over time.

One key challenge of this research into novel interaction techniques with intelligent fabrication machines is that these autonomous systems do not yet readily exist. This research takes a proactive approach in anticipating a near-future abundance of intelligent fabrication machines. However, the realities of today are a lack of existing autonomous CNC machines to develop from. Consequentially, each system I implement must devote significant technical development towards bringing these interactive systems to life. I largely achieve this by implementing interactive systems with levels of *cybernetic intelligence* – where complex behaviors can emerge from lower level systems of feedback and control [Wiener 1961].

Cybernetics often offers a simpler model for artificial life, than artificial intelligence or machine learning. However, it is an effective and valuable strategy for creating artificial entities that exhibit independent and responsive behaviors [Dautenhahn 1999]. I opt for this cybernetic approach since *interaction*, not novel contributions in artificial intelligence or machine learning, is the central concern of this research. Moreover, the interactive systems I implement throughout this body of work demonstrates that this cybernetic approach can foster behaviors with enough autonomy to explore interaction design with intelligent fabrication machines.

This research aims to develop clever solutions to entrenched problems in CAD/CAM, HCI, and HRI. I take a design-oriented approach towards seeking overlooked or underexplored ideas in these technical domains. My work strives to strike a balance between pragmatism and poetry: I tackle the pragmatic problems of today to better tease out the poetic potentials of tomorrow. This tactic has proved useful to hold my work against a high standard of academic rigor and cultural relevance. Moreover, grounding my ambitious inquiries with technical contributions has provided opportunities to distribute my work through both academic publication and artistic exhibition. As a result, the work presented in this dissertation has developed into a research practice capable of engaging wider, non-technical audiences by bringing these future-looking concepts out of the lab and into the collective conscious.

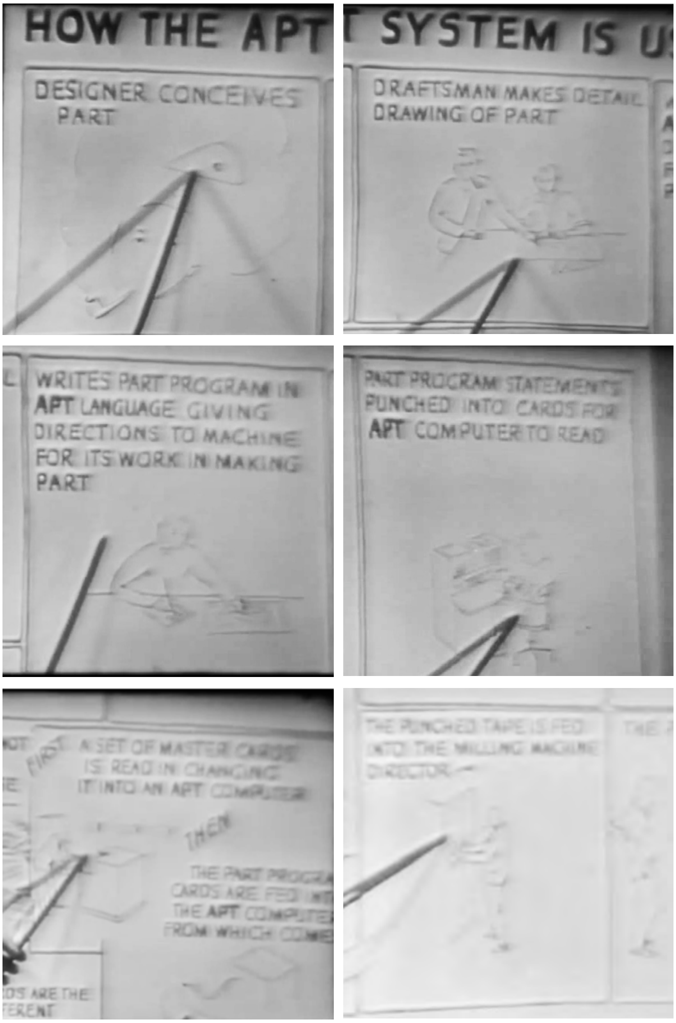

FIGURE 2

Still shots of the MIT Science Reporter television show in 1959 demonstrating the first NC system, the Automatically Programmed Tool (APT). In these stills, Douglas Ross outlines the necessary steps for using APT, which has remained fundamentally unchanged for nearly 60 years.

Background

The Historical Legacy of CNC Machines

In *Forces of Production*, David Noble calls fabrication machines as “the guts of modern industry” for their role in shaping American labor politics in the twentieth century [Noble 1984]. The legacy of these machines is still palpable today, even as trained workers are increasingly supplanted by intelligent automation. In this section, I identify the common methods of control, communication, and interaction for conventional fabrication machines. By better understanding these machines in the abstract, we can more clearly see their advantages, limitations, and future possibilities.

There are a wide variety of fabrication machines used across diverse industrial applications. However they all share the same basic principles in how they operate: the fabrication machine moves its tool tip, its *end-effector*, to a given location (e.g., *an x,y,z point*), with a certain orientation (e.g., *a w,x,y,z rotation*), and then does a simple function (e.g., *to deposit, cut, or grab a material*). No matter the fabrication machine – whether it is a 3-axis printer or a 6-axis robotic arm – they all follow this straightforward pattern of control.

When used to manipulate matter, this relatively simple computational input can result in sophisticated physical assemblies. The aerospace, aeronautical, and automotive industries were eager early adopters of Computer-Aided Design (CAD) and Computer-Aided Manufacturing (CAM) production processes, and were quick to develop innovative applications for fabricating high-precision shapes, such as airfoils [Coons 1963] and complex forms, such as car hoods [Bézier 1986]. The technology behind such systems has seen incremental improvement over the past 50 years, however, there has been little innovation in how these machines operate. A designer today uses a fabrication machines as they might have fifty years ago: they convert a design into digital geometry using a CAD program, they then adjust the geometry to correspond to a fabrication process in a CAM program, and then the digital geometry is converted into machine code and sent to CNC machine (Figure 2).

By contrast, general-purpose robots – machines developed as open-ended platforms for navigating environments – used integrated sensing from their onset, and could see and respond to a changing world. Instead of relying on pre-programmed geometry, these machines combined sensor data with advanced mathematics for kinematic modeling, motion planning, actuation, manipulation as a means of communicating and control. This early emphasis on adaptability and responsiveness is one reason why general-purpose robots have had an explosion of diverse forms and functions over the past half-century, while fabrication machines have remained relatively unchanged.

However, the creators of the first digital design and fabrication machines had loftier ambitions for the impact and evolution of their systems. For example, Ivan Sutherland – whose pioneering system *Sketchpad* laid the foundation for graphical user interfaces (GUIs), Computer-Aided Design (CAD), and Human-Computer Interaction (HCI) in the early 1960s – described his work as a way for "man and a computer to converse rapidly through the medium of line drawings" [Sutherland 1963]. Moreover, Douglas Ross and Steven Coons, authors of one the first Computer-Aided Design/Computer-Aided Manufacturing (CAD/CAM) systems, saw their work as a way to better join "man and machine in an intimate cooperative complex" when conceiving and producing the built environment [Ross 1963]. Here, at the earliest origins of digitally making things, there is an aspirational emphasis on *communication* and *cooperation* with computational machines. These early pioneers built the potential for their systems to become more than tools, and suggested that computational machines could potentially enhance human creativity as collaborators and co-authors.

The Current State of Human-Machine Relations

For the past 150 years, manufacturing technologies have been a driving force in the development of modern society. In the mid-19th century, the transition from human-powered to steam-powered machines led to *mechanized production* processes [Brynjolfsson 2014], and the late-19th century introduced electrically-powered machines for streamlined *mass production* [Fisk 2001], and for the last half-century, digital technologies have optimized machines for *automated production* processes. In each wave of these industrial revolutions, technological innovations in fabrication machinery has been key to increasing human advancement.

Automation – the ability for programmable machines to repetitively produce things without human intervention – has greatly contributed to the unprecedented growth and productivity seen across many industries in the 20th century; including agriculture, medicine, and manufacturing. For example, from 1910 to 2000, the U.S. Bureau of Labor and Statistics reports a 96% decrease in agriculture jobs due in part to the increased productivity of mechanized and automated farming. At the same time, decreasing agricultural employment was matched by significant increases in yields for certain crops during this 90-year period [Wyatt 2005]. Yet despite the many benefits that automation brings – making commodities faster, better, safer, and more economically accessible to wider swatches of the population – automation's promise of prosperity is nearing an inflection point.

Today, the arrival of technologies that displace more jobs than they enable are prompting researchers and policy makers to brace for the next wave of industrial transformation [World Economic Forum 2017]. Rapid advancements in machine learning and artificial intelligence are making fabrication systems smarter and more adaptable than ever. Although these intelligent machines will bring increased efficiency, reliability, and economy to production processes, our direct control and relevance to these autonomous machines will inherently also weaken. With the need for human participation, intervention, and supervision of fabrication machines decreasing, there is a heightened concern for the consequences of a destabilized labor market and wide spread social unrest as societies are segmented into low-skill/low-pay and high-skill/high-pay strata. In moving towards 21st century modes of production, we will need to have more inclusive, human-centered ways for people and fabrications machines to collaborate and co-exist.

REDISCOVERING THE 'COOPERATIVE MAN-MACHINE SYSTEM'

A key challenge to realizing this vision has been the over-reliance on *geometry* as a control system for fabrication machines. While general-purpose robots advanced to more sophisticated methods of sensory input, fabrication machines clung to legacy methods in CAD/CAM to operate. However, geometry alone is not an adequate model for human-machine communication between a designer and a fabrication machine. Geometry works well for drafting and sculpting precise forms from precise materials in precise settings, but it is not flexible enough to accommodate more ambiguous tasks, materials, or environments. For example, the starting and ending points of the design processes are often nebulous: it takes many years of specialized training to be able to consolidate all the formal, functional, material, performative, phenomenal, and aesthetic requirements of a design into well-crafted digital geometry. Moreover, sensory feedback from an environment, a material, or a person can be difficult to encapsulate as static geometry. A designer's intent, a changing physical environment, or a dynamically moving material are valuable streams of external information that can be too unpredictable or ill-defined to represent as static virtual form.

Human-Centered Automation

My research explores how the tools of automation can be repurposed to enhance, expand, and augment human capabilities. Throughout this research, I implement several strategies to appropriate fabrication machines for more human-inclusive purposes. My interfaces strive to make these machine processes more intuitive to use; they work to combine the best-of machine attributes with unique human affordances; they make these processes more adaptive to changing environments; they introduce these machines into domains that have shown resilience towards automation. These strategic layers combine to reconfigure a fabrication machine with a more empathic and nuanced understanding of its environment and human counterpart.

Throughout the history of making things, *computation* and *craft* have been treated as mutually exclusive skill-sets: a software or mechanical engineer might design an artifact, while the machinist or fabricator produces it. Today, these worlds are merging as the software, hardware, and design knowledge for digital design and fabrication becomes increasingly accessible to non-technical audiences [Gershenfeld 2007]. Yet despite the recent rise in democratized and personalized fabrication, many of these processes remain hard to learn *and* hard to master. Significant training and experience is still necessary to acquire the deep software and analog fabrication knowledge needed to incorporate a fabrication machine into a design and production process. Fabrication machines must become easier and more intuitive to use if they are to provide some level of agency for people.

One goal of this research is to lessen the cognitive overhead that this technology can impose on a craftsperson. It uses two primary strategies to simplify digital-physical workflows: it promotes direct, continuous feedback between digital and physical worlds, and it embeds fabrication knowledge directly into software. Furthermore, fabrication machines have an amazing set of affordances that people lack: they have super-human endurance, precision, and speed, and depending on the machine, they could also have super-human strength, reliability, and size. Likewise, people have unique abilities that fabrication machines currently lack. Attributes like adaptability, dexterity, spontaneity, imagination, abstract problem-solving, and critical thinking are all crucial skills for designing and making things (although these attributes may not be exclusive to humans for much longer). If humans and fabrication machines to augment one another, they will need thoughtfully designed interfaces that combine the best of their abilities with the best of our own.

Pairing complimentary attributes between people and machines offers a strategic way to work around current limitations in digital design and fabrication. Fabrication machines are an incredible conduit between digital and analog worlds: they can tangibly manipulate the physical world as their minds natively navigate connected, virtual environments. However, these machines are intended to only operate in highly controlled environments on precisely calibrated materials. Their beneficial affordances – their speed, precision, reliability, and endurance – often come with tradeoffs in flexibility and responsiveness. By contrast, there are important domains not yet supported by computational workflows, as they require far more material and spatial flexibility throughout the design and fabrication process.

Body-centric domains – such as fashion, film, prosthetics, or wearables – are one such industries that have yet to embrace digital design and fabrication into core workflows. Body-centric domains still heavily rely on high-skill analog craft practices for production. Unlike digitally fabricating a boat, car, plane, or product, these processes for the body need to be able to dynamically adjust to moving targets, handle deforming materials, and operate in close proximity to people to make things on or around the body. While this level of precise adaptability is currently beyond the scope of most commercially available tools for digital manufacturing, it is an active area of research across Computer-Aided Design/Computer-Aided Manufacturing (CAD/CAM), Human-Computer Interaction (HCI) and Human-Robot Interaction (HRI).

FIGURE 4

Many segmented steps of CAD/CAM process

Moreover, there are unique human affordances that can be leveraged to strategically bridge these gaps in our current technology. Interfaces that accommodate spatial sensitives and senses, like proxemics and proprioception, can leverage a person's innate understanding of their body in space to foster a spatial awareness for machines cohabitating with people [Hall 1966; Takayama 2009; Ballendat 2010]. Furthermore, interfaces that promote behavioral attributes, such as anthropomorphism and empathy, can tap into our human instinct to project personable characteristics onto inanimate objects. This can be useful in building better legibility behind a machine's intended and actual actions [Breazeal 1999]. Rather than offer direct technical solutions, these strategies improve usability and better acclimate humans and robots to each other's strengths and limitations.

The Future of Human-Machine Relations

This body of research demonstrates how natural gesture, body language, and proxemics can be harnessed to broadcast valuable, nuanced information between a person and a fabrication machine. Fabrication machines offer a fully-matured robotic platform for inventing, testing, and evaluating the future of human-robot relations with large, non-humanoid robots. Although they have working alongside people relatively unchanged for nearly half a century, these robust and reliable machines share similarities with many of the newer species of robots that are just joining us in the built environment. For example, autonomous consumer robots – such as drones or self-driving cars – are also large machines which can potentially harm the people around them. These non-humanoid robots also don't look or act like people, but still need effective ways to communicate with their human counterparts. Overcoming the communication, control, and interaction limitations of today's fabrication machines provides a robust proving ground for testing out many challenges still emerging for these newer robotic platforms.

The fifty-year legacy of human-robot interaction in automation provides a well-charted territory for inventing more meaningful, human-centered connections to powerful, autonomous machines. Better interfaces enable more fluid, intuitive, and safer ways of interacting with these machines that make things, and unlock new possibilities for their unique affordances. A fabrication machine with an intelligent awareness of their context no longer needs to be confined to factory settings: their super-human strength, speed, endurance, precision, and reliability can be redirected towards alternative, non-automation purposes. My research only opens the door to these new potentials. However, I hope the foundation laid here encourages future researcher to continue to instigate, innovate, and invent a spectrum diverse relationships we can have with these powerful, intelligent machines.

Chapter 2

Related Works

This research exists at the intersection of many different disciplines, including design, new media art, human-computer interaction, and human-robot interaction. It builds on rich bodies of work focusing on gestural interfaces, CAD/CAM, computer vision, human-centered design, 3D printing, and industrial robotics. The following works reflect the most influential interactive systems for my research.

FIGURE 1

Researchers in human-computer interaction have been exploring skin as an interface for mobile computing. Skinput, by Chris Harrison (2010), is one such system that detects a person's tactile interactions with digital buttons projected onto the skin.

On-Body Interfaces

On-body interfaces appropriate the body as an always-available, spatio-tangible surface for sensing and displaying digital content. Skin has a set of unique affordances that have enabled a vast array of interaction scenarios to push beyond traditional, screen-based input surfaces. It is a large, tactile, and elastic surface that is always-available to every person. Moreover, *proprioception* – the ability for the body to know where it is in space relative to itself – is exceptionally useful when integrating computing onto the body. This sense enables the body to act as a persistent anchor for the spatial and tactile memory of users, which is important for spatial input [Hinckley 1994]. Since visual feedback is not entirely necessary to effectively interact with skin-based interfaces [Harrison 2012; Lin 2011; Olberding 2013], on-body interfaces also do not require a user's full attention for valid input.

For Mobile Computing

In mobile computing, several techniques have been developed for sensing input on one's skin. This includes sensors worn on the body [Harrison 2010; Yang 2012; Chan 2013; Chen 2013; Ogata 2013; Laput 2014], implanted under the skin [Holz 2012], and embedded in the environment [Gustafson 2011; Harrison 2011]. To display information, projection-based systems can overlay traditional mobile interfaces, including buttons, menus, or games, directly onto the body (Figure 1) [Harrison 2010, 2011]. Projection-based skin interfaces have also been explored as a method to spatially guide the movement of a user [Sodhi 2012], as well as a way of overlaying medical information directly on patients [Ni 2011]. Finally, the social, physical, and psychological implications of on-body input are also emerging as important design criteria for skin-based interfaces [Harrison 2012; Olberding 2013; Weigel 2014].

For Digital Design and Fabrication

Interfaces for on-body digital design and fabrication are a more recent area of exploration, as the body brings an additional set of challenges. To begin, traditional CAD/CAM software do not support digitally crafting wearables; they are most commonly used for architectural, aeronautical, automotive, and product design applications. Moreover, our bodies are highly complex, highly specific physical contexts; 3D modeling tools tend to be empty virtual spaces that ignore physical contexts. This can make digitally designing wearables to have a tailored or ergonomic fit difficult, even for experienced designers. Furthermore, every *body* is different; tailoring a design to the specificities of many bodies types can be a laborious task with current CAD/CAM processes.

On-body fabrication brings even more constraints. Most existing materials for additive manufacturing are harmful to the body: melted plastics, UV-curing resins, and thermosetting resins could scald if they come into direct contact with skin. Moreover, these inflexible materials are not well-suited to deform with the body and are not easily removed from the body. Traditional fabrication machines are also potentially damaging to the body. CNC machines – such as 3D printers, CNC routers, and industrial robots – are designed to operate in highly controlled environments, on completely static stock materials. By contrast, the body is a soft, deformable, and always-moving canvas, which makes fabricating precise digital geometry very difficult. Moreover, existing CNC machines are not intended to be use in close proximity to the body, and can be extremely dangerous if misused.

DIGITAL DESIGN

Despite the inherent challenges of this nascent field, interesting work is being developed in on-body digital design and fabrication. Tactum – discussed in detail in Chapter 4 – was the first interactive system of its kind to explore 3D modeling 3D printed wearables at a directly scale on the body. Since then, other systems have explored digital tools for designing things at a 1:1 scale in complex physical contexts [Ashbrook 2016; Zhu 2016; Kim 2016; Huo 2017], as well as other domain-specific applications for on-body design [Saakes 2016; Nittala 2016].

DIGITAL FABRICATION

On-body digital fabrication is still an emerging field of exploration; however, art and design communities have a strong tradition for crafting designs directly on the body and on proxies for the body using analog techniques. Fashion designers and tailors create bespoke garments on models or mannequins; special effects artists craft prosthetics and props on actors or lifecasts; tattoo artists inscribe their graphic designs onto their client's skin. In many of these domains, however, the techniques for on-body design and fabrication are purely analog. In each of these scenarios, the artifact is customized and hand-crafted on an individual's body by highly trained fabricators.

There are a handful of notable works that have explored on-body digital fabrication in various forms. From the world of fashion, Alexander McQueen's used two industrial robots to spray paint a model's dress as she wore it on the catwalk (*Dress No. 13*, S/S 1999) (Figure 2). Another on-body work featuring an industrial robot is *Tatoué II (2016)* by Appropriate Audiences – an artist collaborative developing CNC machines for tattooing. *Tatoué I*, created in 2014, first modified a desktop 3D printer into a 3-axis computer-controlled tattoo gun. *Tatué II* was a second iteration, which leveraged the additional degrees of freedom of an industrial robot to make a 6-axis tattooing machine (Figure 3). Other explorations include general fabrication techniques for on-body electronics interfaces [Weigel 2015; Kao 2016; Lo 2016].

FIGURE 2

Alexander McQueen's Dress No. 13 featured a pair of industrial robots painting a garment while being worn on the catwalk.

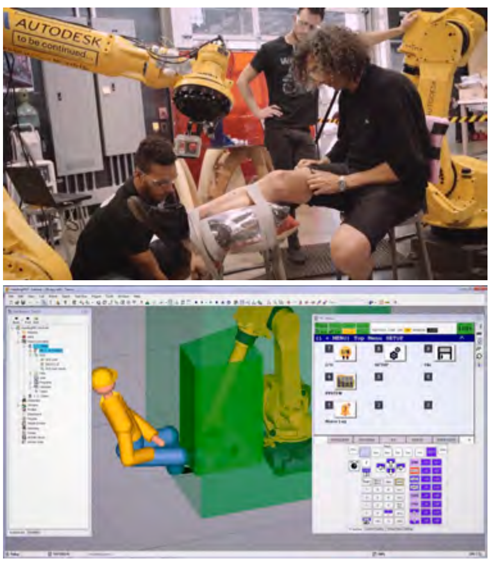

FIGURE 3

Tatoué II, from Appropriate Audiences, transformed an industrial robot into an automated tattoo gun.

The tattooing machines by Appropriate Audiences are particularly valuable examples for on-body fabrication: not only were both systems able to successfully tattoo people, they also bring to light many steep limitations and challenges of adapting existing fabrication machines to work on the body. These systems show that calibration and safety are the biggest obstacles to overcome: in both *Tatoué I* and *Tatoué II*, the body part being tattooed had to be rigidly strapped into place.

While this increased calibration accuracy, it also increased safety risks: in *Tatoué I*, the arm was pinned in between the tattoo gun toolhead and the build platform of the modified 3D printer, and in *Tatoué II*, the leg was pinned between a pedestal and the toolhead. This creates a dangerous pinch-point, where a calibration error, program error, or hardware failure could cause the tattoo toolhead to puncture a body part. *Tatoué II* had an additional calibration challenge: the body part, the digital geometry, and the robot needed to be calibrated to the same coordinate system. 3D scanning the body part helped to synchronize the three independent coordinate systems. Although performed under highly controlled scenarios, under expert supervision, there was a non-trivial risk of bodily injury with these systems.

This literature review reveals that on-body digital design and fabrication is an underexplored area that is ripe for further experimentation. There are very few existing digital tools that cater to the design and fabrication needs of body-centric industries, and consequentially, a wealth of domain spaces have yet to adopt digital craft practices to design and make things for the body. Wearable computing devices – such as watches, smart eyewear, fitness trackers – medical devices – such as braces, splints, prosthetics, or casts – athletic gear – such as cleats, guards, or helmets – and special effects props – such as prosthetics or masks – are all substantial industries lacking digital technologies that integrate into their existing high-skill, analog craft practices.

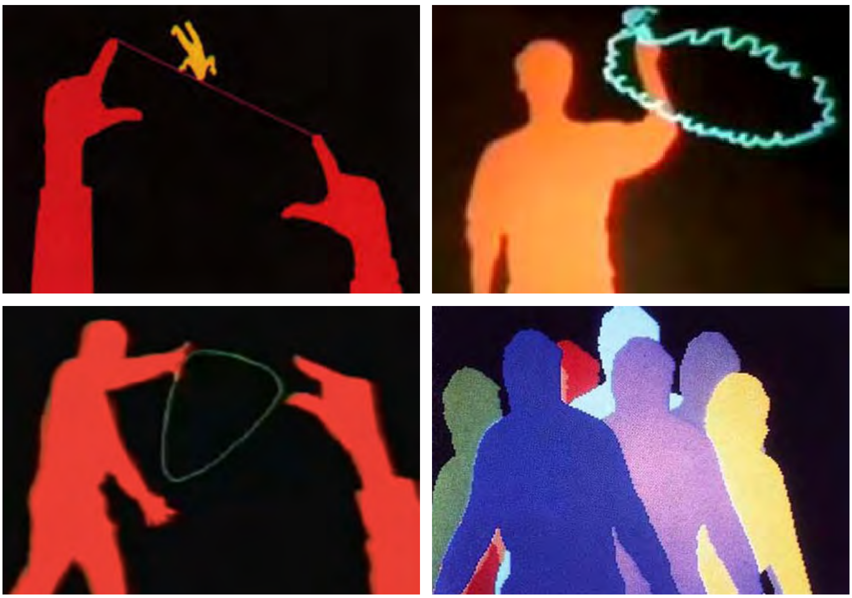

FIGURE 4

Myron Krueger's Videopalace *pioneered gestural user interfaces, and rigorously explored interactions at multiple scales of the body.*

Gestural Interfaces for 3D Modeling

Gestural User Interfaces can directly connect the physical body to a digital world. In facilitating body-centric interactions, gestural UI's bypass traditional human interface devices (HID) – such as the mouse, keyboard, or screen – which act as intermediaries and proxies for engaging with digital content. Moreover, real-world bodily gestures are an interaction model that is more similar to how people actually operate in the analog, physical world. As such, well-designed gestural UI's have the potential to be more intuitive to understand, control, and access than traditional user interfaces.

Early exploration into gestural interfaces was pioneered in the early 1970s by the computer artist, Myron Krueger. His immersive installation *Videoplace* (1972 – 1995) was one of the first interactive systems to illustrate how the body can be used to directly manipulate digital form [Krueger 1977]. Throughout its two-and-a-half decade run, *Videoplace* rigorously explored how gesture could engage with digital content at many scales of the body: from the hand, to the forearm, to the whole body, to many bodies (Figure 4). Moreover, Krueger's foundational work for *Videoplace* gave a vision for many alternative forms of computing, including new media art, touch screens, augmented reality, and virtual reality.

When applied to Computer-Aided Design, gestural UI's strive to anchor the experience of crafting digital form in a spatial, tangible, or physical context. One early application was *Put-That-There*, which connected a user's voice and gestures to manipulate and command projected graphics around a space (Figure 5) [Bolt 1980]. However, more common are gestural modeling systems that use spatial, mid-air interactions to fluidly create or modify 3D forms. These systems are markered or markerless, and often rely on computer vision to detect, track, and classify gestures for 3D modeling. Mid-air gestures can also be used to search databases of pre-existing, highly detailed 3D models [Holz 2011], digitally sculpt virtual forms by manipulating proxy materials, such as

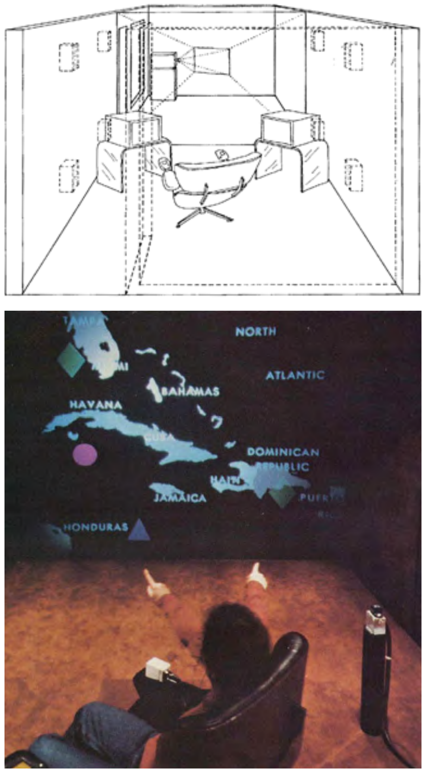

FIGURE 5

Richard Bolt's 1980s system, "Put-That-There", explored the combination of gesture and voice to spatially interact with an immersive computing system.

sponges and clay [Sheng 2005; Smith 2008], and incorporate tangible tools to modify geometry [Llamas 2003].

More recently, researchers in 3D modeling for digital fabrication have begun incorporating physical contexts into the digital modeling environment [Follmer 2010, 2012; Leithinger 2011; Lau 2012]. MixFab, for example, incorporates physical props with gesture-based geometry creation for 3D printing [Weichel 2014]. Bridging analog and virtual contexts helps novices of fabrication based 3D modeling understand a sense of scale in what would otherwise be a scale-less virtual environment.

Despite these benefits, there are steep limitations that have prevented gestural user interfaces from more widespread use. To begin, mid-air interfaces are innately bad at repeating exact bodily movements, as they lack spatial anchors as persistent references [Anderson 2013]. Moreover, these systems have been shown to quickly fatigue a user's arms or body [Hincapié-Ramos 2014]. Additionally, the overall tolerances for tracking a moving body are much higher than that of traditional HIDs, and despite their intuitive use, it can be difficult for such systems to offer fine-grain control when 3D modeling. Consequently, mid-air gestural modeling systems are often limited to sculptural or abstract modeling [Gross 2001; Llamas 2003; Schkolne 2001; Sheng 2006; Zhang 2013] or to simple 3D navigation [Kim 2005; Wang 2011].

This tradeoff between intuitive interaction and precise control is the primary obstacle to incorporating gestural interfaces into CAD/CAM workflows. However, the following chapters presents several interaction techniques that add precision and retain fluid usability for gestural 3D modeling. Moreover, existing gestural modeling systems have yet to take full advantage of certain affordances that are unique the body, such as proxemics – the sense of your body in space – and proprioception – your body's sense of itself. Finally, this research considers the body as a powerful grounding element for gestural interactions – not only for the unique challenges and implementation constraints, but for the deeply personal understanding that connecting digital experiences to our own physiology can bring.

Fabrication-Aware Design & Hybrid Fabrication

Fabrication-aware design and hybrid fabrication systems embed a better awareness of a production process into a digital tool. However, each approach this problem from a different direction: fabrication-aware design embeds the technical expertise of an experienced fabricator into a design-base software system, whereas hybrid fabrication embeds this knowledge into a fabrication-based hardware system.

Fabrication-aware design has been traditionally used for large-scale engineering projects, where multiple material assemblies must adjoin to create complex geometries [Pottmann 2013]. However, the rise of affordable CNC machines (e.g., 3D printers, laser cutters, routers), has made fabrication-aware design an invaluable technique for opening advanced modeling and fabrication to non-expert users. Researchers are looking for new interfaces that make digital design more intuitive, interacting, and engaging for this wider, nontechnical audience [Willis 2011]. Sketch-to-fabrication systems, for example, link sketch-based 2D geometry to additive or subtractive processes in fabrication [Coros 2013; Johnson 2012; Mori 2007; Saul 2011].

Hybrid fabrication, or hybrid craft, integrates digital and analog fabrication techniques to augment traditional craft with digital workflows [Zoran 2013, 2014]. Craft practices such as drawing [Mueller 2015], carving [Zoran 2013], weaving [Zoran 2013], painting [Shilkrot 2015], sculpting [Peng 2015], and fashion design [Wibowo 2012] have been hybridized with digital techniques. Tools developed for hybrid fabrication are often hand-held devices that use mechanical or computational interventions to increase precision and accuracy [Peng 2015, Rivers 2012; Shilkrot 2015; Zoran 2013; Teibrich 2015]. These tools have an awareness of the material they are manipulating and their location in space, and provide visual [24, Rivers 2012; Shilkrot 2015] or tangible [Peng 2015] feedback.

The role of a hybrid fabrication machine during the production process can vary: from passive to neutral to active. With a *passive* approach, digital techniques may be used to print a static formwork that artisans can build upon [Zoran 2013]. A *neutral* approach could involve using digital techniques to guide the user, but not intervene if they deviate [Shilkrot 2015]. With an *active* approach, actuated tool-heads can be used to correct or constrain user actions to match a desired digital model [Peng 2015, Rivers 2012, Zoran 2013].

Fabrication Aware Design and Hybrid Fabrication systems share common limitations. Although these systems can make CAD/CAM processes more intuitive for a user, there are tradeoffs in agency and control between a user and a final physical output. Similar criticisms have been raised for hybrid fabrication processes: *Being the Machine* reflects on the power relationship between a user and hybrid fabrication machine [Devendorf 2015]; *Hybrid Artisans* examines the value added and value lost to traditional craft practices [Zoran 2014]. Finally, fabrication-aware design tools are typically limited to designing around simple physical contexts, and hybrid fabrication systems primarily rely on static build volumes and canvases.

Fabrication-aware design and hybrid fabrication are both are useful techniques for integrating the best affordances of digital and analog making. This body of research present techniques for balancing ease-of-use with purposeful authorship in hybrid tools for digital making. Chapter 5, in particular, presents design considerations for creating such systems, and pushes the complexity of forms that can digitally designed and fabricated on and around. Moreover, it emphasizes techniques for designing and fabricating digital form on complex, moving, and deformable surfaces.

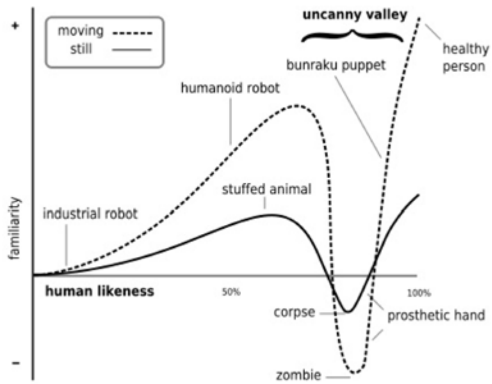

FIGURE 6

Mori's well-known diagram mapping the "Uncanny Valley" describes the cognitive dissonance felt when interacting with almost-human machines.

FIGURE 7

Star Wars's C-3PO (left) and R2-D2 (right) represent two opposite strategies for the mechanical and interaction design of intelligent robots.

Robots with Personality

Pop culture and science fiction have solidified personable robots as a part of our collective conscience. Feature films, such as *2001: A Space Odyssey*, *Star Wars*, and *Wall-E*, all illustrate a future where humans and machines co-exist alongside one another. Arguably, however, the most engaging of these fictional robots are those that don't emulate humans. Instead, these non-humanoid robots use alternative means of communication to make a deep, personable connection to an audience.

The design of these robots tap into an ambiguous space between two psychological phenomena: *pareidolia* – an instinctual impulse to see meaningful features, such as faces, animals, and objects, in inanimate things – and the *uncanny valley* – a cognitive dissonance caused by identifying imposter features in human-like things (Figure 6) [Mori 1970]. Within this region of believability, a person will innately suspend their disbelief, and their minds will actively fill in the missing elements that would otherwise make a thing seem alive. Moreover, this lower-definition interpretation may even solicit a more active engagement from an audience than would a truer-to-life version [McLuhan 1964]. However, once beyond this believability threshold, the relationship switches: a person will instinctually notice the ways in which a thing is un-human [Moore 2012].

Perhaps the most elucidating example of this tension in robotics comes from the robot duo made famous through the *Star Wars* film franchise: R2-D2 and C-3PO (Figure 7). The design of R2-D2 could be affectionately described as no more than a trashcan on wheels, and it only communicates to the world through an intelligible language of beeps and boops. By contrast, C-3PO is a humanoid robot that shares many of our same bodily features, facial features, and idiosyncrasies. However, throughout the film series, R2-D2 seems far more personable and lifelike, and is believably cast as the plucky side-kick of the main protagonist. C-3PO, for all its human-like android features, seems far more rigid, artificial – and in the end – more robotic.

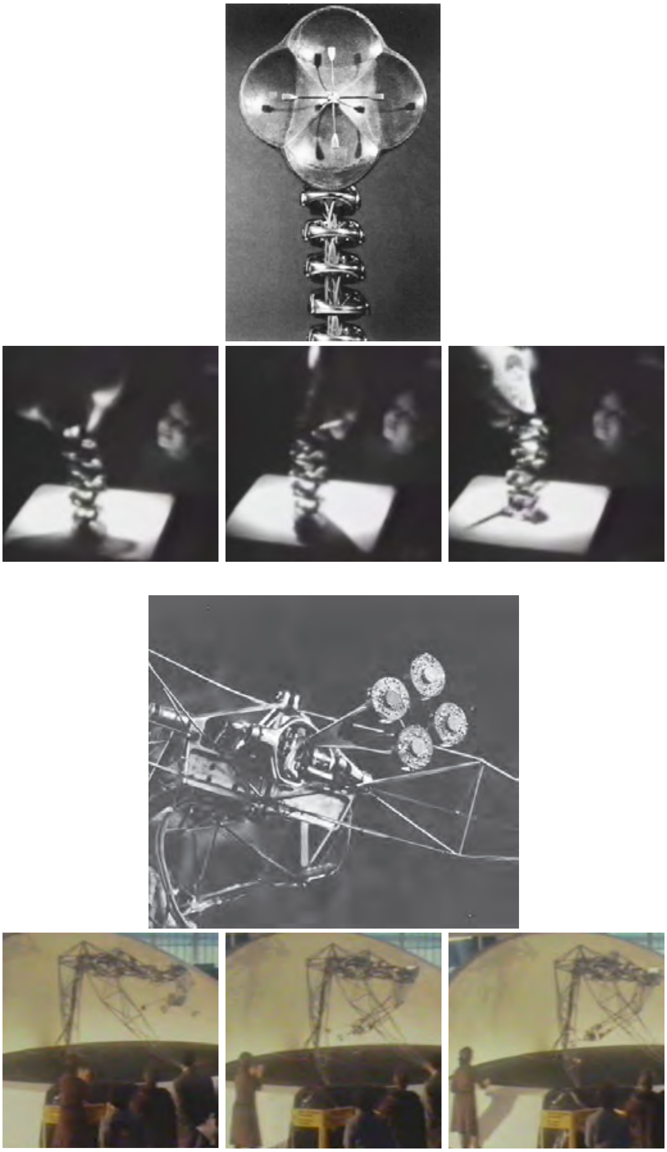

FIGURE 8

Ihnatowicz's S.A.M. – Sound Activated Mobile *(top) and* Senster *(bottom) are one of the first robots to be designed as unique mechanical creatures.*

In Human Robot Interaction (HRI), the metric of a robot's intelligence or personability is often based on its likeness to people. While this is a prominent philosophical framework, there are also other useful design metaphors worth exploring. One of the most notable alternatives comes from the work of Edward Ihnatowicz (1926-1988) – a pioneering roboticist and robotic artist who saw robots as creatures, not things. His two seminal pieces, S.A.M. (*Sound Activated Mobile*) (1968) and *Senster* (1970), are often categorized as *cybernetic sculptures* [Reichardt 1969], however, these mechanical creatures had equally sophisticated sensing, actuation, and control systems as other contemporary robots.

S.A.M and *Senster* were both imaginative, original creations that didn't follow the existing aesthetic norms of how a robot should look or act (Figure 8). Ihnatowicz's mechanical design, when paired with clever audio-based sensing and custom hydraulic actuation, brought these mechanical creatures to life as seemingly independent cognitive systems [Reffen 1984]. Both S.A.M and *Senster* used audio as an external stimulus, and relied on a mechanical system (in the case of S.A.M.) or a computer-controlled system (in the case of *Senster*) to reorient their bodies towards the direction of the audible sound. Although Ihnatowicz did not set out for his robots to directly emulate animal movements, they were nonetheless evocative of familiar, lifelike kinematics [Zivanovic 2005].

A more recent example that follows this trajectory of inquiry includes Golan Levin's *Double-Taker (Snout)* (2009), a giant, Henson-esque robot that checks out passersby (Figure 9). *Double-Taker* is a standard industrial robot, dressed in a googly-eyed costume that is reminiscent of "an enormous inchworm or elephant's trunk" [Levin 2008]. Like S.A.M. or *Senster* before it, *Double-Taker* uses its unique body to directly engage with visitors. However, it uses computer vision, rather than audio, to sense its world: when it detects a person of interest, it reorients its supersized googly-eye towards them. This not only suggests to passersby that *Double-Taker* is intelligently aware of its environment, but it also prompts a recursive dance: the robot moves surprisingly

FIGURE 9

Golan Levin's Double-Taker (Snout) *is one of the first examples of body-language being explored on an industrial robot.*

to a person's presence, which then causes a person moves surprisingly to the robot's presence.

Both Ihnatowicz and Levin orchestrate a fluid call and response between their robots and visitors to facilitate legible, empathic connections. However, one important technical and conceptual distinction between Ihnatowicz's robotics creatures and Levin's *Double-Taker* is the relationship between sensing and actuation. Much like our own sensory organs, the sensors in *S.A.M* and *Senster* are attached to their presumed faces: their microphones pick up a signal and their faces reorient towards that signal. In short, their ability to sense their environment is limited by their bodies. However, *Double-Taker* is not bound by its body: how it sees the world and how it acts in the world are profoundly decoupled. *Double-Taker* may look at you with its googly-eye, however it does not see you with it; instead, it uses a camera from a different, fixed vantage point. On one hand, this separation may reflect a 40-year shift from mainframe to ubiquitous paradigms in computing. However, it is also illustrative of one of the most unique and underexplored affordances of interactive robotics: with their brains detached from to their bodies, these machines can act as a tangible bridge between virtual and physical worlds.

Chapter 3

Reverb
Designing with Autonomous Geometry

Rather than begin working directly with machines, my research into human-centered interfaces initially began by focusing on intelligent, autonomous geometry. Geometry is the underlying communication structure for engaging with fabrication machines. However, today there are still steep limitations to the intelligence, usability, and interactivity of virtual geometry for CAD/CAM.

This chapter illustrates a series of techniques for overcoming such limitations, and demonstrates how to transform geometry from static virtual form into an intelligent, animate collaborator in the design process. It presents *Reverb*: a context-aware 3D modeling environment that lets you design ready-to-print wearables around a 3D scan of the body. *Reverb* is built around principles of human-centered design: it encapsulates information about a user, a virtual canvas, and fabrication process into a contextually aware design tool, *Reverb* enables a designer to craft digital things around precise physical contexts, which embeds a level of ergonomic intelligence into its virtual environment. In offloading this cognitive overhead, *Reverb* strives to show how intelligent, autonomous geometry can make such systems easier and more intuitive to use for designers.


The Body as Digital Context

Reverb explores how to craft intricate, digital geometries around complex physical contexts, however, it focuses most extensively on designing around the body.

The body is a highly complex, dynamic, and personal physical context that brings its own unique affordances and challenges. At its most pragmatic, designing for the body is a great use case for 3D printing, as the build volumes most 3D printers are the right size and scale for the body. Moreover, digital fabrication processes – like 3D printing – are well-suited for the added variability required to customize a design to each unique body [Gershenfeld 2007; Piroozfar 2013]. Furthermore, there is currently a lack of commercially available CAD tools that offer solutions for the body. There are research prototypes that focus on simulating the drape

FIGURE 1

Reverb is a context-aware 3D modeling environment that lets you design ready-to-print wearables around your own body.

or movement of cloth [Volino 2005], generalizing models of human anatomy [Wang 2009], interactive 3D modeling [Umetani 2011], deconstructing 3D designs for 2D fabrication [Mori 2007], and tangible interfaces for garment design and fabrication [Wibowo 2012; Yamashita 2013]. However, these systems are not widely available, and they only focus on cloth-based garment design; they do not account for the design of wearables made from rigid or semi-rigid 3D printed materials.

CAD tools for architectural, aeronautical, automotive, and product design applications are widely available. However, a building, a boat, a car, or a gadget has inherently different design challenges than the highly complex, highly specific context of the body. To begin, the empty void of a virtual

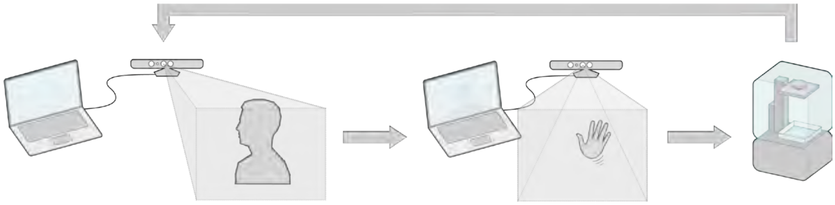

FIGURE 2

Reverb's cyclical workflow digitizes a physical context via 3D scanning (left), enables 3D modeling around the scanned context (middle), and incorporates the form into the phyical context via 3D printing (right)

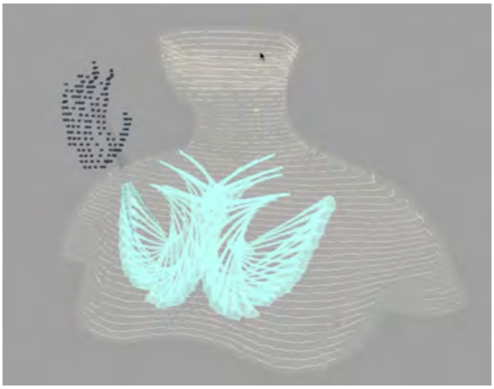

FIGURE 3

Reverb enables a designer to drape intelligent, fabrication-aware geometry around a 3D scan of their own body.

modeling environment is ill-suited for the non-Euclidean contours of the body. Furthermore, current design tools do not allow digital form to be crafted around an actual physical context; they, instead, situate a digital design within some arbitrary, context-less Cartesian grid. Additionally, since every *body* is different, there is a wide range of formal variation across body types. This can make creating variations of a design an extremely laborious task. Moreover, the body can move: it can squish, stretch, bend, and shake. Digital tools for designing wearables – such as accessories, prosthetics, medical devices, athletic gear, or footwear – need to account for these dynamic, external forces.

Reverb

Reverb captures a designer's mid-air hand gestures to craft intricate digital geometries that can be immediately printed and worn on the body. It uses a three-phase workflow to facilitate the capture, design, and fabrication of a wearable artifact around the body (Figure 2). In the capture phase, *Reverb* uses a depth sensor to 3D scan the body as a persistent physical context in its virtual environment. The design phase uses this same sensor to continuously track the mid-air hand motions of the designer. These hand motions attract the attention of an autonomous virtual creature inside *Reverb*, enabling the designer to drape this squid-like form through space and time around the 3D scanned context (Figure 3). In the fabrication phase, the designer can immediately export a ready-to-print mesh from *Reverb*, once a desired design has been generated. The final physical artifact results in complex lattice structures that both conform to and expand on the body.

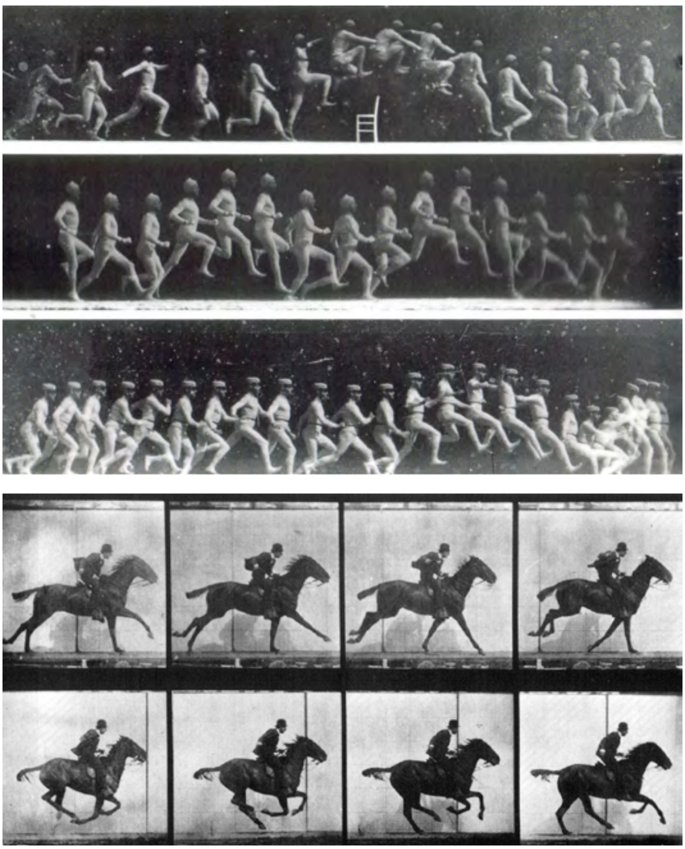

FIGURE 4

The 3D modeling technique is Reverb was initially inspired by chronophotographic experiments of Marey (top) and Muybridge (bottom).

Implementation

Hardware Configuration

The workstation for *Reverb* is comprised of a first-generation Kinect mounted on a tripod to face a designer. An auxiliary screen is placed near the designer to visualize *Reverb*'s virtual environment; however, it cannot occupy the tracking region of the Kinect. The effective tracking region is 100cm x 100cm x 100cm, with an approximate 3-5mm working resolution. A designer sits behind the rear edge of this tracking region, and reaching their hand inside to interact with the virtual environment. The system was implemented using the Java programming language at 15-30 FPS, and uses following the open-source libraries: processing for graphics, OpenKinect for depth imaging, and toxiclibs for physics simulation.

3D Scanning

A designer steps into the tracking region to 3D scan themselves during the capture phase of *Reverb*. The software allows for multiple scans to be captured, since the depth camera used in *Reverb* does not allow for full 360° scanning. The designer can then manually orient and position of each captured point cloud using a GUI panel to create a fully three-dimensional scan. Pre-existing scans can also be imported into *Reverb* as .stl or .csv files.

Hand Tracking

Hand tracking in *Reverb* is a naïve and straightforward implementation. To detect a 'hand', the system tracks the closest 100 points of the point cloud streaming from the depth camera. With the designer is seated behind the tracking region, it is a reasonable assumption that the point cloud detected by the depth camera will be a hand. The centroid of this point set is used as the focal point for a designer's interactions with the virtual environment. Using multiple hands has the added effect of rapidly moving this focal point through the tracking region.

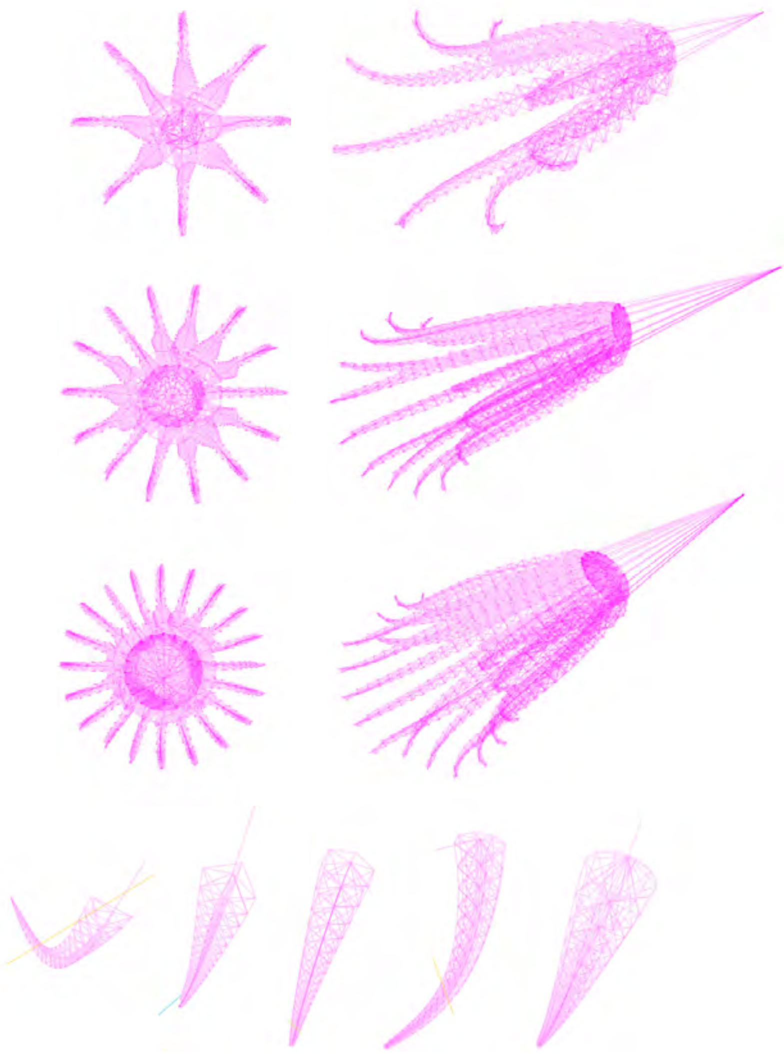

FIGURE 5

Reverb's virtual agent is parameterized to change the length and joint locations within the legs, the profile of the legs, and the number of legs. This, in effect, adds formal variation to designs while retaining the same overall aesthetic.

3D Modeling

Reverb takes an ambitious approach to building intelligent, empathic digital design environments: it aims to enable designers to intuitively 3D model with interactive, sentient, fabrication-aware geometry. To achieve this, *Reverb* synthesizes a unique combination of 3D modeling techniques into one digital design tool, including *chronomorphologic*, *parametric*, *physics-based*, and *agent-based modeling*.

CHRONOMORPHOLOGIC MODELING

Reverb uses a single, interactive geometric module to build intricate three-dimensional forms through a technique I developed called *chronomorphologic modeling*: during the design phase, the system records a copy of module at given time intervals as it moves through the virtual environment. Recording can be triggered manually, using a keypress, or automatically, when the virtual creature nears the bounding box of the 3D scan. This novel modeling technique was inspired by the early *chronophotographic* experiments of photographers, such as Étienne-Jules Marey and Eadweard Muybridge. These early pioneers developed several methods for capturing the dynamic motions of quickly moving humans, animals, and objects in two-dimensional photographs (Figure 4). *Reverb* adapts this 2D motion capture to 3D modeling, allowing a designer to extrude the movement of the module through space and time.

PARAMETRIC MODELING

Reverb also integrates traditional parametric modeling functions to allow for added variation and experimentation across a single aesthetic family (Figure 5). For example, local parameters – such as the number of legs, length of legs, or joint elasticity of the module – and global parameters – such as the viscosity, time, or gravity in the virtual environment – are all adjustable through a traditional GUI panel.

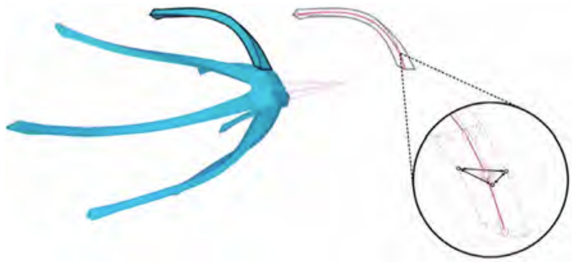

FIGURE 6

The virtual agent's spring-based particle mesh is implemented to retain the printable quality of manipulated geometry. Additionally, particles in the spine of its leg are weighted to create joint-like behaviors.

FIGURE 7

The shape and form of the virtual agent impact its fluidity and density as a base geometric module.

PHYSICS-BASED MODELING

To 3D print a 3D model, digital geometry must satisfy three conditions: it must have a minimum thickness, it must be made of closed-meshes, and these meshes cannot have self-intersecting faces. Rather than fix invalid geometry to meet these conditions, our design tools could avoid the problem altogether by only allowing properly formatted geometry to be created.

In *Reverb*, this is achieved by constructing user-manipulated geometry from a spring-mass model simulated in a virtual physics environment. The squid-like base module in *Reverb* is built on the physics-based spring-skeleton: each edge of its mesh faces is connected by simulated springs, and spines of repelling particles that keep its body inflated at all times (Figure 6). This elastic skeleton adds a pseudo-proprioceptive quality to the geometry: it preserves the module's fabrication-aware properties, no matter how the module is manipulated or agitated by a designer's hand. Additionally, digital geometry is first initialized as a closed mesh with a minimum thickness parameter. This minimum thickness is maintained, and self-intersecting faces are prevented, by inflating the spring-mass model with repelling particles. This physics-based approach ensures that geometry inside *Reverb* can have fluid and dynamic movement, while innately adhering to these strict fabrication-aware properties (Figure 7).

AGENT-BASED MODELING

Early experiments with *Reverb* revealed that a designer's hand movements were not varied enough to offer significantly different motion paths for its base module. This resulted in repeatedly generating overall forms that were too self-similar. To counteract this homogeneity, I programmed a level of autonomy into the base module, so it could move and act on its own. This internal logic and agency for navigating the virtual world is based a variation on Craig Reynold's seminal boids algorithm for agent simulation [Reynolds 1987].

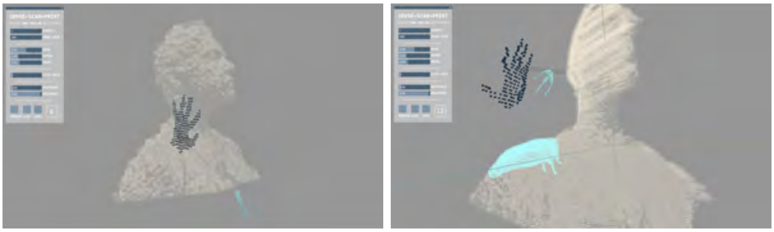

FIGURE 8

A designer's hand interacts and guides Reverb's autonomous geometry around the digitized model of their body.

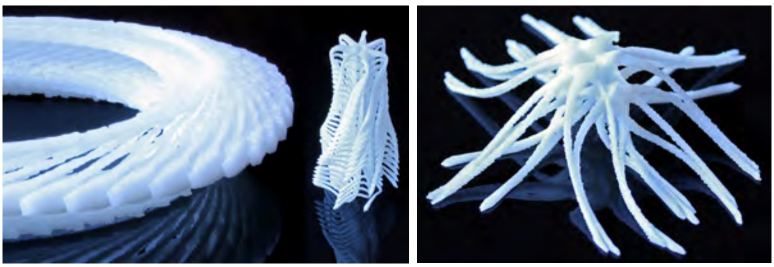

FIGURE 9

The first 3D prints created with Reverb verified that fabrication-aware properties were preserved throughout the design and printing process.

FIGURE 10

3D modeling around the wrist tested the maneuverability and control for designing wearables around a highly three-dimensional physical context.

Adding these autonomous behaviors impacted the relationship between the designer and design tool in *Reverb*: it transformed a geometric module that *looked like* a creature into an *actual* virtual creature. With this second order level of control, the designer only indirectly influences the behavior of an independent virtual entity: their hand controls the point-of-interest that the virtual creature could autonomously navigating towards. These behaviors – when combined with the other modeling techniques layered into *Reverb* – generated enough variation for designs to engage the body in novel ways. Moreover, it illustrates an aspirational, collaborative relationship between the designer and a sentient design tool (Figure 8).

Physical Artifacts

Several wearable artifacts designed with *Reverb* were fabricated through Fused Deposition Modeling (FDM) and Selective Laser Sintering (SLS) printing techniques. These prints were useful for testing the fidelity of the fabrication-aware properties of *Reverb*'s digital geometry. Early experiments with Fused Deposition Modeling (FDM) printers and ABS plastic tested the fabrication-aware properties of the composite chronomorphologic model. This set of 3D prints verified that a physics-based modeling strategy could allow a geometry to innately retain its printablity while being manipulated by a designer: they were immediately exported from *Reverb* and sent to a 3D printer without any boolean or post-processing operations (Figure 9). Later experiments focused on the usability of *Reverb*'s mid-air gestural interface. The forearm was used as a highly three-dimensional and complex virtual canvas to design around. While symmetrical and asymmetrical wearable forms were successfully designed and fabricated for the forearm, repeatability and precision during modeling was difficult (Figure 10). A final set of necklaces were designed around 3D scans of the bust and printed via SLS printing (Figure 11).

FIGURE 11

A final set of necklaces were 3D printed via SLS printing using a soft yellow elastomer and a rigid vwhite nylon.

Discussion

A key contribution of *Reverb* was the newfound ability for designers to intuitively craft precise digital forms around complex physical contexts. To achieve this, the system had to overcome several challenges with mid-air gestural interfaces.

One hallmark limitation of mid-air gestural interfaces is their lack of spatial precision and repeatability (Holz 2011; Fraser 2013). In early experiments with *Reverb*, it proved difficult for a designer's hand motions to guide geometry *around* the 3D scan instead of *through* it. Brilliant mesh algorithms currently exist for detecting and fixing three-dimensional intersections, however these can be both complex and computationally expensive. By contrast, *Reverb* circumvents the issue altogether: rather than fix invalid geometry to meet a given condition, it only allows valid geometry to be created. Moreover, point clouds of a scanned context in *Reverb* are embedded with an active particle system that repels user-manipulated geometry when it comes too close. Similar to how the spring-skeleton of *Reverb*'s base module prevents self-intersections, this integration prevents the same geometry from intersecting with 3D scanned contexts.

This lazy technique of avoiding, rather than solving, a mathematical problem helps tame a designer's expressive, yet imprecise, hand gestures for draping intricate digital forms around a physical context. However, there are additional benefits, beyond circumventing this long-standing limitation. Most pragmatically, bypassing resource intensive mesh-checking algorithms allows *Reverb* to consistently run at interactive framerates. This directly impacts the experience of *Reverb*, as it facilitates intuitive, fluid feedback between a designer's actions and the resulting geometry. Moreover, bypassing the post-processing of invalid geometry helps to prevent disjointed interactions between the designer and the interface. Finally, working around a 3D scanned context that actively influences user-generated geometry can embed a level of ergonomic intelligence directly into the underlying structure of a digital design environment.

This intelligence is most clearly demonstrated in the set of 3D printed necklaces created with *Reverb* (Figure 12). The bust presents a challenging virtual canvas to design for: a designer must guide a geometry on, over, and around highly specific curvatures of the body. However, *Reverb*'s context-aware approach produces strangely anatomical artifacts that also push the absolute tolerances of a fabrication technology. In just a few seconds of working with *Reverb*, a designer can create exquisite forms that conform to and deviate from the body: delicately embracing the body, balancing on the shoulders, or gently resting on the nape of the neck. Engraining ergonomic considerations for how the geometry meets the body enables digital geometry created through *Reverb* to simultaneously reference and expand on its physical form.

Reverb led to a series of design strategies that would become formative to this body of research – including interaction techniques for taming mid-air hand gestures for precision 3d modeling, and 3D modeling techniques for persevering fabrication-aware properties of interactive geometry. Furthermore, *Reverb* marked my first foray into discovering the affordances of working with autonomous digital creatures. These early experiments provided valuable, hands-on experience for developing autonomous behaviors for inanimate things, and helped foster an intuition for defining the right balance of agency and control in the relationship between a person and an independent, artificial counterpart. Finally, the work produced through *Reverb* provided an unexpected opportunity to engage with a range of interdisciplinary communities, including architecture, design, technology, and fashion: architectural publications validated its technical and conceptual merit [Gannon 2014; Gannon 2017]; the project was covered by various tech-oriented media outlets; and the physical prints were exhibited at a 3D printed fashion event for New York Fashion Week in 2013.




FIGURE 12

These printed necklaces complement and expand on the body in a way that indicative of the ergonomic intelligence embedded into Reverb.

While this external validation was certainly comforting, the true value of these experiences was in seeing how a highly technical and narrow idea could bridge design, technology, and culture, and become relevant to a wider, non-technical audience. This led to me to seek opportunities for appropriating tools and techniques from a wide variety of disciplines, and was influential in how I approached future work.

However, *Reverb* was not without its limitations and shortcomings. To begin, the squid-like base geometry in *Reverb* was hard-coded: although the module was parameterized for formal variation, every artifact designed in the system clearly belonged to the same aesthetic family. An internal design tool for building new modules – constructing novel anatomies and ascribing different autonomous behaviors – would have been very useful for exploring different formal starting points when 3D modeling with *Reverb*. Its greatest limitation, however, came from its interface.

Reverb successfully harnesses a designer's expressive, yet imprecise, hand gestures to drape intricate digital forms that fit precisely around a physical body. However, its configuration was not ideal for facilitating fluid and immersive interactions with a virtual environment. Its mid-air interface was challenging to use for extended periods of time and was prone to arm fatigue [Hincapié-Ramos 2014]. Moreover, its auxiliary screen did not provide adequate visual feedback while 3D modeling, and complicated the use of a mouse for navigating through the virtual environment. Finally, real-world hand gestures captured through *Reverb* have a palpable disconnect from the 3D modeling actions inside the virtual environment. While Reverb enabled a designer to reach into a virtual environment, the ideal configuration would bring the virtual environment out of the computer and into the physical world.

Conclusion

This chapter demonstrates how to enhance standard digital geometry with an enhanced understanding of its design and fabrication contexts. Reverb achieves this by combining several 3D modeling and computer vision techniques to embody digital design environments with the knowledge of an experienced fabricator. Moreover, the physical artifacts produced through Reverb illustrate that ergonomic intelligence can also be engrained in a virtual environment when we use a digitized body as a responsive base canvas. Additionally, supplying Reverb's modular geometry with autonomous behaviors shows promise as a method for pushing beyond the limits of a designer's imagination: trading some level of agency and control to an autonomous system can introduce systemic variations and foster unanticipated results. Finally, Reverb provides several successful design strategies for creating interactive, intelligent, fabrication-aware geometry. In the following chapter, I explore how this intelligent geometry can come out of the computer and occupy dynamically changing physical environments.

Chapter 4

Tactum
Skin-Centric, In-Situ Design

A central theme of my research examines ways to facilitate embodied interactions between a person, a technology, and an environment. *Reverb* worked well to embody the knowledge of an experienced fabricator into fluid, intuitive software. However, its interface was not an ideal configuration for immersive interaction: *Reverb*'s mid-air hand gestures were physically and conceptually disjointed from the activities happening in its virtual environment.

My next system, *Tactum*, expands on the scope of *Reverb* by bringing intelligent, autonomous geometry out of the computer and into the physical world. *Tactum* is a fabrication-aware design system that captures a user's skin-centric gestures for 3D modeling directly on the body. Like *Reverb*, digital designs generated through *Tactum* can also be immediately 3D printed and worn back on the body. However, *Tactum* greatly improves on the set of gestures captured for 3D modeling. This novel system presents a unique experience where a designer can digitally craft a design in-situ, at a 1:1 scale on the body.

The process of developing *Tactum* – the first on-body modeling system of its kind – revealed many design considerations for creating other interfaces that use skin-based input for gestural 3D modeling-to-fabrication. This chapter shares insights into ideal configurations and use cases for such systems, and discusses interaction techniques for three different modes of skin-centric 3D modeling. Additionally, it presents several printed, wearable artifacts designed with *Tactum*, and shares a set of observations from design professionals.

Skin-Centric Design

Tactum investigates *skin* as a more relevant, context-aware interface for designing 3D printed wearables. This large, always-available surface enables intuitive, tactile interactions [Ogata 2013; Weigel 2014] and can even be reliably accessed without visual feedback [Lin 2011; Ogata 2013; Weigel 2014]. However, skin-based input has yet to be explored outside the domain of mobile computing. In this chapter, we examine skin as an interactive input surface for gestural 3D modeling-to-fabrication systems. Gesture-based interfaces for 3D modeling offer a number of unique affordances that can empower non-expert users to participate in digital design. They facilitate the expressive creation of digital geometry, while requiring little prerequisite skill for most interactions [Gross 2001; Smith 2008]. However, their intuitive use comes at a cost: it is difficult for these systems to enable both

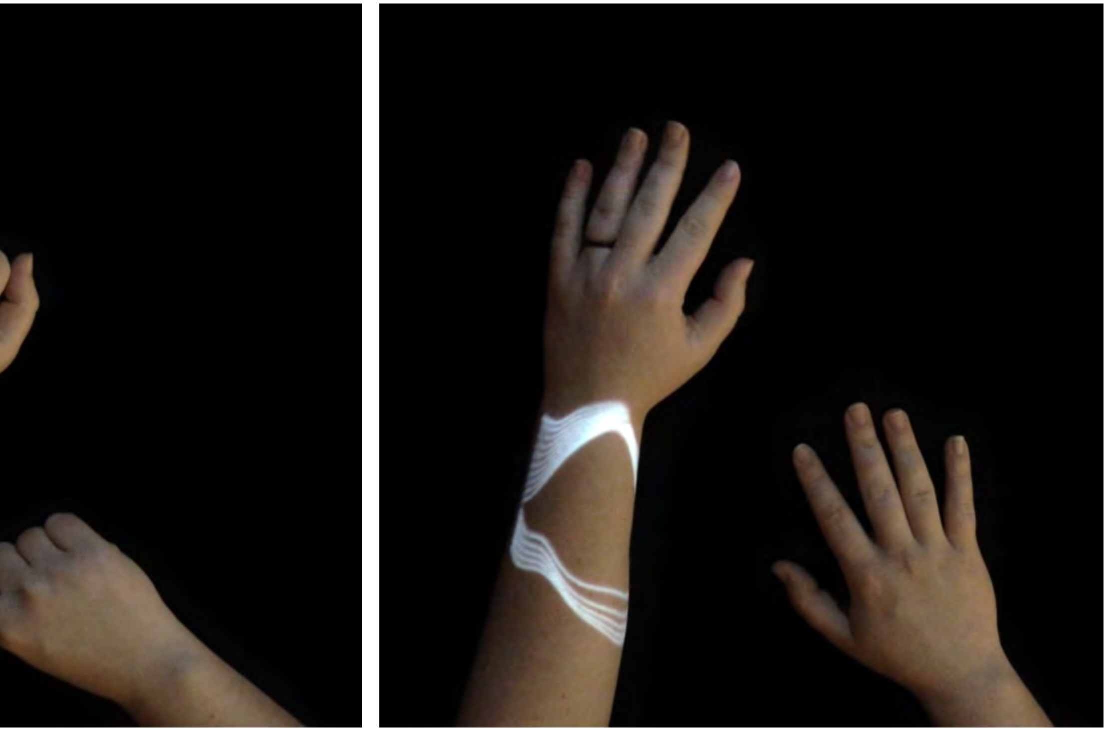

FIGURE 1

Tactum is a fabrication-aware design system that captures a user's skin-centric gestures for 3D modeling directly on the body.

high precision control and expressive form generation. As a result, the digital geometry generated is often limited to abstract or sculptural forms [Schkolne 2001; Llamas 2003; Kim 2005; Sheng 2006; Zhang 2013].

Skin, as both the input surface and base canvas for digital design, can enable non-expert users to intuitively create precise forms around highly complex physical contexts: our own bodies. Furthermore, as new forms of 3D printing and digital fabrication are reaching wider, non-technical audiences, there is a potential for users to design and fabricate personalized products [Gershenfeld 2005]. In some cases, such personalization may relate to a user's own body – such as jewelry, braces, and other wearable devices. As such, design and fabrication workflows that utilize skin as an input platform are worthy of further exploration.

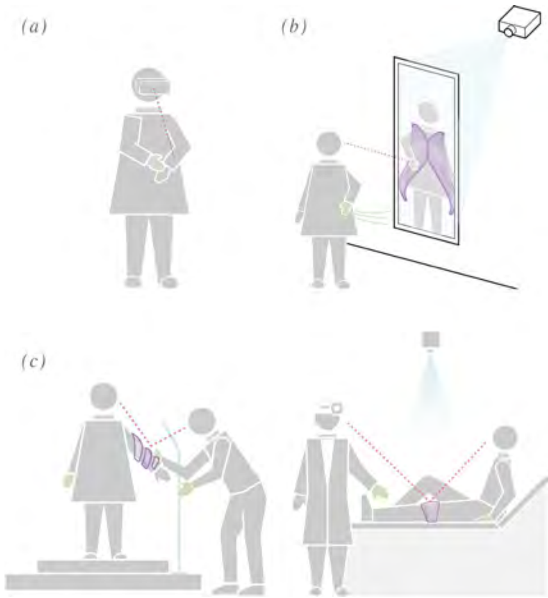

FIGURE 2

The point-of-view of the user or users determine the physical configuration of a skin-centric design tool: (a) firstperson systems have a direct line of sight to the user's canvas area, (b) second-person systems have partially occluded canvas areas, and (c) third-person systems enable multiple users to participate in design.

Design Considerations

This section details design considerations around the appropriate content, configuration, and input and output of gestural modeling-to-fabrication tools.

Possible Content

Skin is an appropriate input surface for design tools intended for on-body artifacts. The non-Euclidean nature of our anatomy can make the design of wearable objects unintuitive in conventional CAD environments. However, appropriating skin as a starting canvas could help embed ergonomic principles into the foundations of a design. We therefore see three domain spaces that could specifically benefit from skin-centric interfaces for digital design.

First, fashion items, like garments, shoes, accessories, and jewelry can be adapted to skin-based input and 3D printing. Second, wearable computing devices, such as watches, smart eyewear, and fitness trackers, can also be customized or personalized using skin-centric design and fabrication tools. Third, medical devices – such as braces, splints, or casts – can be designed, customized, personalized, and fabricated through these systems for at-home rehabilitation.

Possible System Configurations

Three important aspects of a skin-centric design tool's system configuration are the user's point-of-view, the location of the canvas area, and the number of users interacting with the system.

POINT-OF-VIEW

Figure 2 illustrates three possible point-of-view configurations. In first-person systems, the point-of-view of the modeler is directed at their own body [Harrison 2010]. This method is appropriate when there is a clean line of sight between the user and the canvas area. The forearms, hands, and upper thighs are likely locations for this system (Figure 2a). Second-person systems have parts of the desired canvas area occluded from the user's line of sight. As a result, these systems should provide representations of the body through an auxiliary display [Zhang 2013]. The face, neck, bust, back, or full body are likely locations for second-person systems (Figure 2b). Third-person systems have the canvas area located on a person other than the modeler [Ni 2011]. This method is appropriate when multiple users collaborate on a single design (Figure 2c).

SINGLE VERSUS MULTI-USER SYSTEMS

Recent research has surveyed the propriety of touch for different locations of body-based interfaces [Harrison 2011; Ogata 2013; Weigel 2014]. While the social acceptance of touch may not be applicable to first- or second-person systems, it becomes an important factor when designing multi-user interfaces. Skin is useful as a collaborative modeling platform in scenarios where professionals work with non-experts. This scenario could occur with a doctor working directly on a patient [Ni 2011], a fashion designer working directly on a model, or an engineer working with a consumer. It is important to give each party agency in the design process, although one user may have more influence over the final design than another. Moreover, additional instrumentation, such as touching with a stylus instead of the hand, can be introduced in locations where touch is necessary, but socially inappropriate or awkward.

INPUT

Hardware options for detecting touch input on the body have increased in recent years, however not all techniques are applicable skin-centric 3D modeling interfaces. Skin-based input for gestural 3D modeling must be able to detect both tactile input from on-body interactions and spatial input from near-body interactions. Moreover, sensor readings must be translatable to local and global Cartesian coordinate systems. Therefore, many of the hardware solutions that infer touch based solely on disruptions in electric or acoustic signal may provide insufficient information for 3D modeling purposes [e.g., Lin 2011; Harrison 2012; Chan 2013; Chen 2013; Mujibiya 2013]. However, pairing these devices with optical sensors, such as RGB, IR, or depth cameras, can provide robust information for 3D modeling with skin-based gestures.

OUTPUT

Visual output for skin-centric design tools can exist both on and off the user's body. As mentioned previously, one valuable affordance of skin as an input surface is that it facilitates the user's tactile and spatial memory of the body. Therefore, conventional off-body displays can effectively aid a user's bodily interactions. However, there are also several output devices that can provide direct visual overlays. Mobile or embedded projectors (Figure 2c) can provide robust visual feedback, especially when mapping two-dimensional forms onto parts of the body [Gustafson 2011; Harrison 2012]. For depth-rich three-dimensional designs, augmented reality devices, such as translucent screens (Figure 2c) or head-mounted displays (Figure 2a), may be better suited to overlay 3D visuals on the body. However, these devices are still limited by the canvas areas a user can directly see.

Spatial Landmarks

In addition to touch input, individual variations of skin, such as freckles, veins, and tattoos, can provide spatial landmarks to anchor a user's spatial and tactile memory [Gustafson 2011]. If integrated into the modeling workflow, these landmarks can provide a persistent reference to skin-based interactions as the designer works at 1:1 scale with their body.

FIGURE 3

The initial workstation for Tactum was comprised of a Microsoft Kinect mounted above the work area to track user input, and Microsoft Surface tablet mounted within the work area to display visual output.

Tactum

These design considerations for skin-centric design were further examined through the development of *Tactum*, an augmented modeling tool that lets you design 3D printed wearables directly on your body. This system transforms a depth camera into a touch sensor to detect tactile interactions with your own skin. An above-mounted projector displays an animate, interactive 3D model mapped onto the body: a designer can simply *touch*, *poke*, *rub*, or *pinch* the geometry projected onto their arm to customize ready-to-print, ready-to-wear forms. Once a desired form is generated, a designer closes their hand to export their design for 3D printing.

Implementation

Hardware Configuration

The hardware configuration of *Tactum* was implemented through two iterations. The first iteration was comprised of a first-generation Kinect mounted ~800mm above a workstation, with a Microsoft Surface Pro 3 mounted on a desk. The desk is also used as a base surface to place the arms (Figure 3). In this setup, the effective tracking region is 90cm x 65cm, with an approximate 3-5mm working resolution. The second iteration of *Tactum*'s hardware swapped the above-mounted Kinect for a Leap Motion hand sensor. We found that the Leap Motion sensor data had less noise than the Kinect, and provided more reliable robust skeletal tracking during touch interactions. This iteration also integrated a pico projector above the workstation that bypassed the auxiliary screen altogether, and projected visual feedback directly onto the body.

The system was implemented using the Java programming language, and uses following the open-source libraries: processing for graphics, opencv image processing, and toxiclibs for physics simulation.

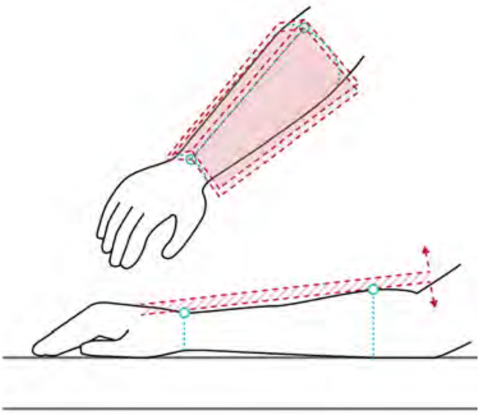

FIGURE 4

Image segmentation and touch detection are generated from the anatomy of the forearm. The medial axes of the arm create the 3D plane for segmentation. An offset from the 3D plane defines the 3D touch region of the forearm.

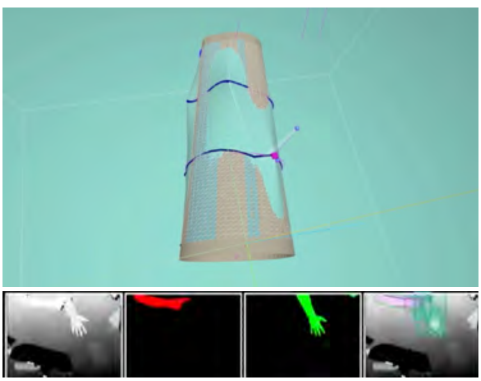

FIGURE 5

Real-time visual output shows the 3D modeling environment, and a computer vision panel with the raw depth image, modeling hand segmentation, canvas arm segmentation, and hand and gesture tracking. points update the modeling hand.

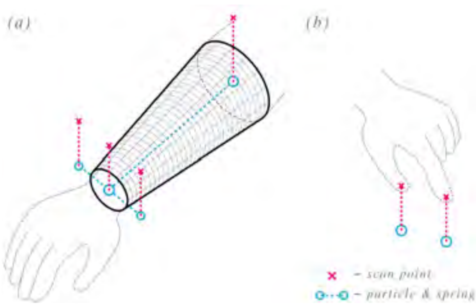

FIGURE 6

Virtual geometry is updated by live scan data from the depth camera. (a) elbow, wrist, and wrist axis points update the canvas arm, and (b) index and thumb points update the modeling hand.

Arm Tracking

In the first iteration of *Tactum*, its depth camera detects a user's forearm as they sit at the workstation. The contours of the arm are then processed to define anatomical regions (e.g., the elbow, wrist, hand, fingers). With these regions defined, we simplify the forearm's point cloud into two 3D planes. These 3D planes dynamically segment the depth image into a modeling hand and canvas arm mask. Everything above these planes becomes a part of the modeling hand mask, and everything between the planes and a given maximum distance becomes the canvas arm mask (Figure 4). The second iteration of *Tactum* leverages the Leap Motion's built in skeletal tracking to detect the arms, hands, and fingers of a designer.

Touch and Gesture Detection

To reliably detect skin-centric interactions between the modeling hand and canvas arm, our first-generation hardware configuration adapts the image processing techniques in [Hinkley 2011] to the three-dimensional geometry of the forearm: the dynamic 3D planes used to segment the modeling hand and canvas arm masks are offset by a minimal distance (~20mm in our implementation). When a finger of the modeling hand enters the space between the offset and original planes, we know a touch has occurred. The second iteration with built-in skeletal tracking follows a similar method to detect touches. This enables skin-centric gestures that will be described later.

Real-Time Display

FIRST-GENERATION HARDWARE CONFIGURATION

In the first iteration of *Tactum*, flat panel display above the work area is used to provide users with real-time visual output of the arm and finger tracking, within the context of a 3D modeling environment (Figure 5). Once the canvas arm is detected and processed, the 3D modeling environment is generated on the auxiliary display. While the system's depth camera can continuously 3D scan the user's forearm, this data was found to be too noisy and incomplete to provide

FIGURE 7

The second iteration of Tactum *integrates interactive projection mapping to visualize 3D models directly on the body.*

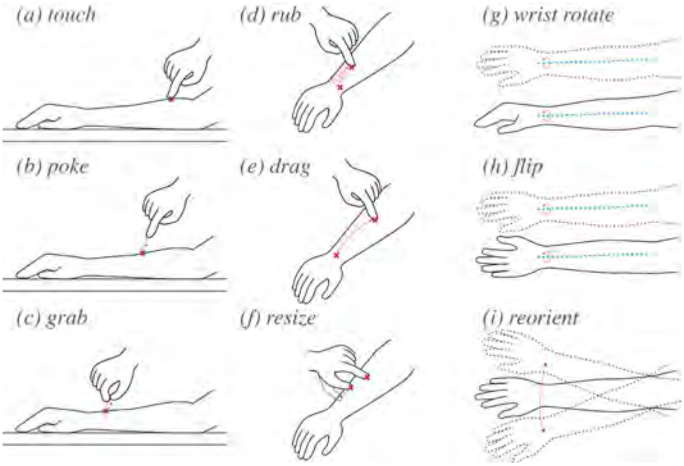

FIGURE 8

Natural gesture set. (a) touch; (b) poke: a single tap on forearm; (c) grab: pinch thumb and index, and touch forearm; (d) rub: touch and drag repeatedly; (e) drag: touch and move; (f) resize: touch and move thumb and index; (g) wrist rotate: rotating hand about the wrist; (h) flip: flipping the forearm; (i) reorient: moving the forearm.

a useful virtual reference in the modeling environment. Instead, we built an accurate representation of the forearm using a simulated spring system in the modeling environment. Live scan data from the depth camera then updates select particles in the spring system with a small number of 3D points from the elbow, wrist, and hand (Figure 6a). The spring system both dampens depth noise from the streaming points and facilitates smooth motion for the virtual forearm. A similar strategy is used to visualize the modeling hand in the virtual environment: a 3D scan point from each finger is rigged to a heavier particle with a spring (Figure 6b). This dampened particle is visualized as the user's virtual finger.

SECOND-GENERATION HARDWARE CONFIGURATION

The second configuration of *Tactum* substitutes an above-mounted projector for the auxiliary display. In this configuration, the 3D model of user-manipulated geometry is projected and mapped directly onto the body (Figure 7). This provided more tactile, immersive visual feedback for a designer, as their gestures manipulated the on-body 3D model. Tactile interactions with the body are streamed to a separate CAD backend that updates the 3D model. Although the CAD backend is not visible to the designer, its updated 3D model is mapped to the body and projected back onto the arm.

Skin-Centric Gestures for Design

Gestures within *Tactum* are designed to be as natural as possible: as you touch, poke, or pinch your skin, the projected geometry responds as dynamic feedback. Although these gestures are intuitive and expressive, they are also imprecise. Their minimum tolerance is around 20mm (the approximate size of a fingertip). While this is adequate when designing some things, it is inadequate for designing wearables around precise, existing objects. This system implements both natural and symbolic gestures based on the tracking of the modeling hand and canvas arm. Figure 8 summarizes the gesture set of our system. Gestures performed with one or two fingers of the modeling hand are

FIGURE 9

Skin-centric gestures incorporating landmarks are generated through an initial calibration step: a user selects landmarks by touching desired points on the forearm, then records a gesture by drawing a path to each point.

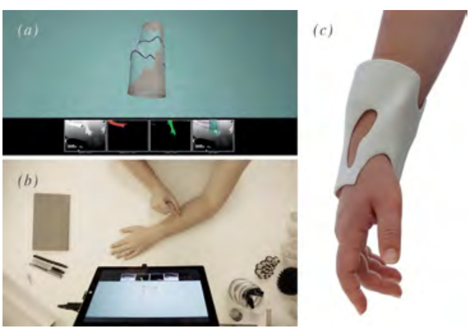

FIGURE 10

An armlet is designed using the parametric manipulation mode. (a) skin input updates the arm's base geometry, and simulated rubber bands can be pulled and deformed by the user; (b) the model can be customized to the user; (c) the resulting physical artifact on the user.

used for design operations. These include touch, poke, grab, rub, drag, and pinch. Additional gestures performed with canvas arm are used for 3D navigation. These include flip, orientation of arm, and rotation of wrist.

The system also includes skin-centric symbolic gestures by incorporating landmarks on a user's forearm. For example, a user can touch their middle knuckle to export a design for fabrication, or they can touch a set of freckles in a particular order to run an application-specific command. An initial calibration step generates landmark gestures: the user first selects landmarks by touching the desired points of their forearm, then records a gesture by dragging their finger to each landmark (Figure 9).

Modeling Modes and Workflow

Tactum illustrates how skin-centric input can be used as a gestural 3D modeling-to-fabrication system through three distinct modeling modes: direct, parametric, and generative. These modes are adapted from the primary 3D modeling techniques of conventional CAD environments. Although bodily interactions and geometric manipulations may vary, the goal of each mode is the same: balance high precision control with expressive gestures, and to ensure fabricated digital designs have an ergonomic fit to the geometry of the forearm.

Direct Manipulation

With direct manipulation mode, the interactions between modeling hand and canvas arm directly transform the base geometry built from the canvas arm. Gestures such as grab, drag, wrist rotate, flip, and reorient are used to modify an underlying mesh structure.

In our prototype, we use direct manipulation to design and fabricate an armlet for the forearm (Figure 10). When the user's forearm is detected, the system generates a malleable digital surface from four control edges built from the virtual forearm (Figure 10a). These control edges are digitally simulated rubber bands that are manipulated by the user's

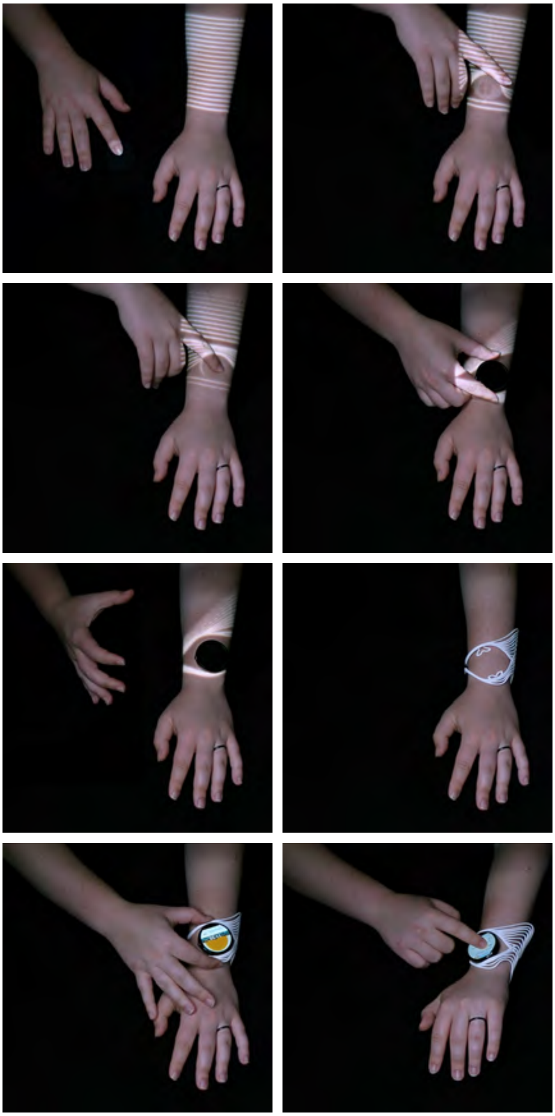

FIGURE 11

Tactum's *on-body interface facilitates 3D modeling around pre-existing physical object. In this demo, the designer uses a smart watch as a 1:1 reference for precisely crafting a custom watch band.*

touch (Figure 10b). Using a grab gesture, the user touches a corresponding spot on their physical arm and pulls the virtual point off their body. An elasticity threshold built into the control edge releases the deformed edge once breached, thereby updating the underlying surface.

This lower level manipulation of geometry is most similar to conventional gestural modeling systems, and it therefore brings similar limitations: while it allows high formal variation, it is difficult for users to have precise control over the final form. To compensate for this lack of control, we ensure the digital design will have an ergonomic fit to the body by subtracting the volume of the virtual forearm from the final manipulated surface. The geometry can then be exported for fabrication, and the 3D printed artifact can be placed back on the body of the user (Figure 10c).

Parametric Manipulation

Within parametric mode, the user's gestures interact with open parameters of a pre-designed digital form. This mode allows a base design generated by an expert to be manipulated by a non-expert. Gestures such as touch, poke, resize, flip, and reorient are used to manipulate and stimulate an interactive parametric model.

To illustrate this modeling mode, I used the second-generation, projection-based version of *Tactum* to customize the design of a watch band for a Moto360 smartwatch (Figure 11). As the user's canvas arm is detected, the system attaches an array on particle strings to their virtual forearm, while projecting the geometry onto their physical body. Like *Reverb*, this animate, intelligent geometry is embedded with fabrication-aware properties: the design will be ready-to-print no matter how it is manipulated by a user. Once this animate, fabrication model is attached to the forearm, the designer can use their opposite hand to interact with it using skin-centric gestures, like touch, rub, and poke.

Interactions with the actual body are sent to a CAD backend for geometry processing. In this example, the designer begins by touching their wrist to set the position of the watch face. They can then make a pinching gesture to resize

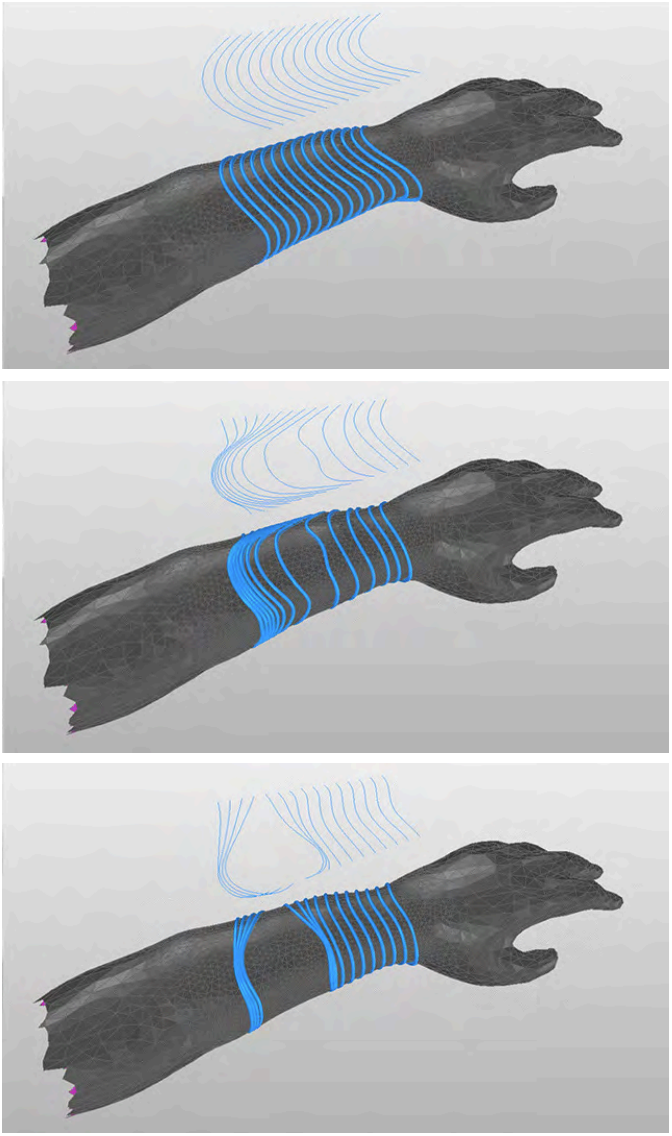

FIGURE 12

A CAD backend projects Tactum's interactive geometry around a 3D scan of the body, out of view of the user.

the opening for the watch face. Since the designer is working at a 1:1 scale on their body, they can also use the actual watch face as a physical reference when setting the size of the opening. Next, they then pinch-and-touch to define the overall form of the watch band. The user then closes their canvas arm hand when a desired form is finalized to export the on-body design for 3D printing.

In this example, skin gestures enable the designer to manipulate some – but not all – of the parameters of the watchband model. Open parameters, such as the position and orientation of the watch face, are able to be updated by a designer's touch. However, parameters that require more intricacy or precision are closed to skin gestures. For example, the clips to hold the watch face and the clasp to close the band onto the arm must have exact measurements and tolerances for the fit, function, and fabrication of the watch band. The detail required here goes beyond the capabilities of gestural 3D modeling, so the exact geometries for the clips and clasp of the smartwatch are instead topologically defined within the band's parametric model. The clips and clasp, while dependent on the overall watch band geometry, cannot be directly modified by any gestures. Once the user has finalized a design, *Tactum*'s CAD backend places and generates these precise topologically defined geometries into the final, ready-to-print 3D model (Figure 12).

While designs generated through parametric manipulation may have limited variation in form, they facilitate both expressive gesture from the user and a high level of quality control through the pre-designed module. Furthermore, this modeling mode illustrates techniques for 3D modeling around existing objects, and for balancing intuitive gesture with precision control for gestural 3D modeling. In the current implementation, this module must be pre-programmed, however future development could create an interface for user-defined parametric modules. This mode may be particularly useful for personalizing consumer products where high tolerance precision is required.

FIGURE 13

An arm brace designed using the generative manipulation mode. (a) skin input updates a touch heatmap of the user's forearm; (b) a design is generated using virtual agents; (c) resulting geometry is processed for fabrication; (d) the resulting physical artifact on the user.

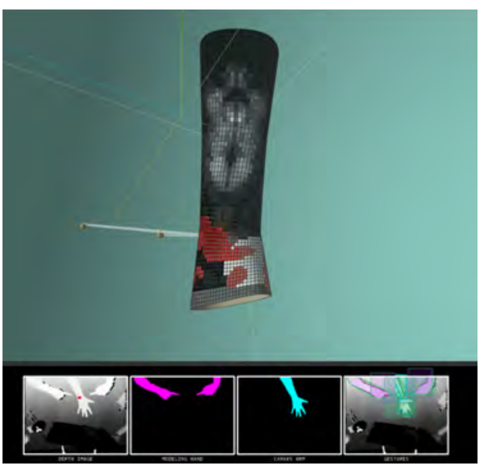

FIGURE 14

In multi-user mode, users can have a different level of impact over the design. Here, one user adds suggestions (red), and the other user implement the design (white).

Generative Manipulation

With generative manipulation, a user's gestures manipulate the underlying abstractions that guide the behavioral properties of an expert defined design. Gestures such as touch, rub, and drag are used so a user can guide how a design is regenerated on the forearm.

Our prototype enables a user to personalize the structure and support of an arm brace (Figure 13). When the user's forearm is detected, the system generates a touch heatmap around the forearm, wrist, and hand of the user. Areas where the user rubs become brighter to indicate repetitive touch along the heatmap. Once the user achieves a desired pattern for the heatmap, they use a landmark gesture to generate a digital design, and the system then deploys a predesigned generative algorithm that processes the touch heatmap. In our implementation, we program hundreds of virtual agents to seek out bright areas and avoid dark areas of the user-defined heatmap. As they move across the three-dimensional surface, the trails left behind each vehicle are used as the digital geometry. In effect, areas where the user touches more add extra support for the arm brace, and areas where the user touches less receive less support. Once the simulation has finished, the geometry is exported and post-processed for 3D printing.

For this example, we have also implemented a multi-user mode, where a client and professional user design a brace together. In this scenario, touch interactions from the client are given less weight than the professional: client touches appear in red, whereas professional touches appear in white on the touch heatmap (Figure 14). The client can therefore indicate where they would like brace supports to be generated, but the professional controls the actual location and density of the generated design.

Generative manipulation strikes a balance between the formal variation of direct manipulation and the precision and control of parametric manipulation. It enables the user to directly influence highly complex geometry, but can also ensure quality control over design and fabrication parameters. This mode may be particularly useful in

FIGURE 15

Artifacts developed with Tactum were verified by printing designs on FDM, SLS, and SLA 3D printers.

scenarios where a high amount of user agency is desired in the design process of a complex artifact, such as in personalized medical devices or prosthetics.

Physical Artifacts

Tactum has been used to create a series of physical artifacts around the forearm (Figure 15). The fidelity of printed geometry created through Tactum was tested by printing artifacts on 3 kinds of 3D printers, using a variety of materials: artifacts from the direct manipulation example were printed with ABS and PLA plastic on a Fuse-Deposition Modeling (FDM) printer; examples from parametric manipulation were printed from resin on an SLA printer and nylon on an SLS printer; and the generative manipulation example was printed with nylon and rubber on an SLS printer.

Initial User Observations

Since *Tactum* is intended for both expert and non-expert users, 10 participants with varied backgrounds and experience levels in 3D modeling and 3D printing were invited to participate in an observation and feedback session. Three participants were design professionals, two were professional artists, two were design students with engineering backgrounds, and three were design students with architecture backgrounds. The 3D modeling and fabrication experience of each participant varied from complete novice to seasoned expert.

Procedure

To begin each session, the setup, sensing, and different skin-centric interactions for gestural 3D modeling were explained to participants. Then each of the three modeling modes were demonstrated, highlighting their differences. After each demonstration, participants used the system and were guided through the gestural modeling process. Participants also tried on sample artifacts made through our system. Below we summarize the key observations and comments collected from these sessions.

Participant Feedback

In general, participants gave positive feedback about *Tactum*. They were all able to create or manipulate digital models with the three modeling modes, although most participants took longer to create a satisfactory design with direct modeling than with parametric or generative. When prompted for thoughts on skin as an input surface for 3D modeling, participants were enthusiastic about bring touch and tactility to the digital design processes. One participant made a comparison to a haptic mouse:

> *"I've used a Phantom [haptic mouse] before for 3D modeling, and it's kind of cool, but it still feels like poking something with a stick instead of actually touching it." (p5)*

They also appreciated how a digital design would inherently be scaled to fit. One novice to 3D modeling noted:

> *"It makes sense that if you're designing something to wear on your body, you should be able to literally design it on your body." (p1)*

However, we received mixed reactions on the usefulness of landmark gestures. Most participants liked the novelty, but didn't see it as an improvement over clicking or pressing a button:

> *"I guess it's really only useful when you have to keep both arms highly engaged in modeling." (p4)*

> *"It's neat, but I could also just press a key to export my file." (p2)*

In general, participants with little experience in 3D modeling appreciated the direct and parametric modeling modes most, and found them "engaging" and "empowering". Those with experience in 3D modeling and printing tended to favor the parametric and generative modes. When discussing the generative mode, one experienced modeler noted:

> *"I have no idea how I would recreate that in the 3D modeling software I'm used to." (p3)*

However, the experienced modelers showed concern with how the model definitions would be created. One participant suggested:

> *"I don't really want to have to learn another parametric modeling software [...] it would be great if you could just import the parameters from Grasshopper or Maya or something." (p9)*

All the participants with experience in 3D printing liked that they didn't have to think about formatting the digital geometry for fabrication. One beginner noted,

> *"It's great that it takes care of that for you [...] I always get nervous about it right before I go to print something I modeled." (p6)*

Overall these observation sessions show that both expert and non-expert users were able to design wearable artifacts using skin-centric gestures. They were also eager for additional forms to manipulate in parametric mode, and offered suggestions for other algorithms to incorporate into generative mode.

Discussion and Future Work

The initial results and observations with *Tactum* were encouraging. Despite the preliminary nature of these evaluations, our observational study provides insights into the feasibility of skin input for fabrication-aware design. However, our work only begins to explore what could be a wide design space around skin-centric input for design. As our system develops further, more thorough evaluations will be necessary to fully understand issues around ease-of-use and ergonomic fit for printed designs. This section reflects on lessons gathered from our implementations and observation sessions, in addition to topics for future work.

Both expert and non-expert participants showed enthusiasm for what can be produced through *Tactum*, however experts were concerned with how parametric and generative content could be created. For this research prototype to engage design professionals, back-end content creation would need a friendlier, non-programing interface. This back-end interface could be designed as a stand-alone, modular design system, or it could integrate workflow pipelines from existing design-to-fabrication software.

This implementation relies on solely on skin-based input to push the envelope of all user interactions. However, as noted by some participants, traditional desktop input devices could also be used in combination with the skin-based gestures we developed. Integrating voice commands could be an effective means for hands-free communication with the gestural modeling system.

Additionally, this implementation relies on a single depth camera to sense skin-input on the forearm. While this one sensor did allow for many gestures to be recognized, issues of image resolution and occlusion were inherent limitations that we had to negotiate. To reconcile the low-resolution and noisy data of the Kinect with the high resolution and fidelity of the user-manipulated 3D geometry, we chose to visualize the forearm and hands in the modeling environment as simplified representations. However, directly visualizing the skin and hand in the virtual modeling environment may help strengthen the connection between a user's on-body interactions and the manipulation of digital geometry. Using additional cameras or sensors to capture and map a user's actual skin texture and hand details onto user-manipulated geometry could help bridge bodily interactions and geometry manipulation.

Relying a single sensor also limited the fidelity of *Tactum*'s tracking. While this system allowed a user's free arm movement to manipulate digital geometry, only canvas arm gestures (e.g., flip, rotate, reorient) directly manipulate geometry through free movement. In two-handed interactions, our tracking and recognition strategy requires the canvas arm be stationary to adequately detect gestures from the modeling hand. Integrating additional modes of sensing into skin-centric gestural modeling system could allow for more dynamic interaction between both arms. It could also enable better detection of skin-specific gestures.

There are several output modalities that are appropriate for skin-centric design systems. In our implementation, we elected to use an auxiliary display, as it provides consistent visual representation without suffering any possible occlusions. However, direct visual overlay onto the user's forearm (through head-mounted displays, projection mapping, or see-through displays) could also provide a strong connection between bodily interactions and geometry manipulation. Projection, in particular, could be a useful means of visual overlay. Future work could investigate how projection could also display three-dimensional volumetric geometry that comes off the forearm, such as the geometry created in our direct and parametric manipulation examples. Although our system implements basic multi-user interactions, there are more ways for skin to be a collaborative input surface for design. For example, we show how two users can collaborate on one forearm. But an alternative scenario would be for two users to remotely work on one design on their own forearms. This scenario can increase in complication when design teams work on a single design.

Finally, *Tactum* focuses only on gestural modeling on the forearm. However, there are other parts of the body that are appropriate for skin-centric design tools: the face and head can become an interactive canvas for designing eyewear, masks, or apparel; the shoulders and neck can be used to design jewelry or medical braces; full body systems can be used to design costumes or garments; legs and feet can be used for shoes, medical braces, or casts. Moreover, skin-centric tools for deformable body parts, such as the joints of the neck, back, elbow, knee, or ankle, could combine complex mechanical design with intuitive customization.

Conclusion

This chapter demonstrates how skin can be used as an input surface for gestural 3D modeling-to-fabrication systems that enable non-expert users to create highly complex and expressive designs for the body. Moreover, the observation session with design professionals and students using *Tactum* indicates the high potential for skin-centric gestural modeling to be a collaborative platform for experts and non-experts. Furthermore, the design considerations for future skin-centric systems outlined in this chapter identify an array of different configurations, contexts, and domain spaces for digital on-body design. Additionally, the three separate modeling modes developed through *Tactum* illustrate tangible manipulation techniques for balancing geometric precision and expressive control in such systems. While advances in hardware have increased accessibility to 3D printing, software interfaces have yet to provide increased agency in who can use this technology. However, skin can act as the interface that bridges digital and physical contexts, and can better enable experts and non-experts alike to participate with this technology. In the following chapter, I explore the possibilities of not just digitally designing, but also digitally fabricating wearables on the body.

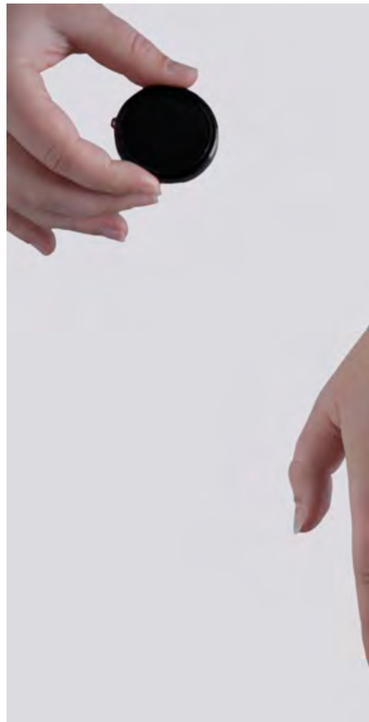

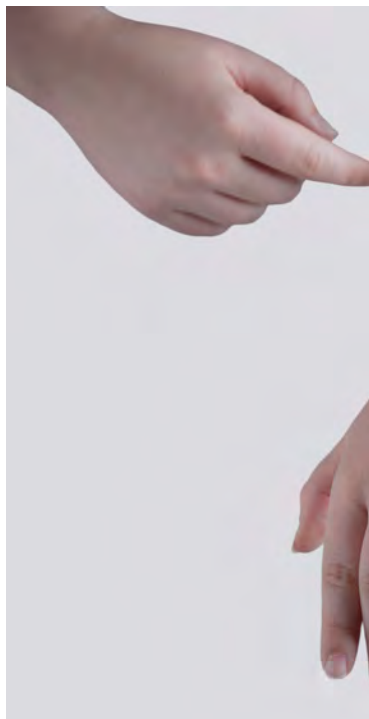




Chapter 5

ExoSkin
On-Body Fabrication

So far, we have discussed *Reverb* and *Tactum*, which demonstrate how autonomous, fabrication-aware geometry can be brought out of the computer to interact with people in dynamic physical environments. These systems promote human-centered interactions with design environments by encapsulating excessive technical information about 3D modeling, 3D printing, and ergonomic design into fluid, intuitive design tools. However, they also present a strict bottle neck at the point of fabrication. Although *Reverb* and *Tactum* both enable a person to craft intricate, precise digital designs in seconds, it can take dozens of hours before they can hold the physical artifact in their hand. Moreover, the immersive fluidity of the digital design process is not mirrored during the fabrication process, since a designer must essentially freeze their design into a fixed, static geometry before exporting to a black-box machine.

My next system, *ExoSkin*, works to increase continuity between on-body design and production processes. This context-aware fabrication system facilitates a range of new workflows for digitally designing *and* printing wearables directly on the body (Figure 1). *ExoSkin* does not try to replace or supplant traditional analog craft practices for the body. Instead, it provides a series of analog-digital workflows that enhance these high-skill, analog techniques with digital tools; providing guidance and feedback, and supporting digital operations such as importing, exporting, and archiving designs.

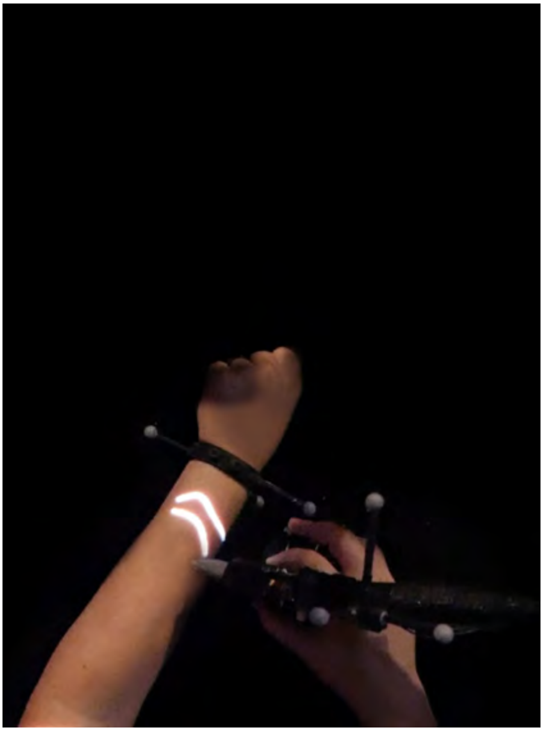
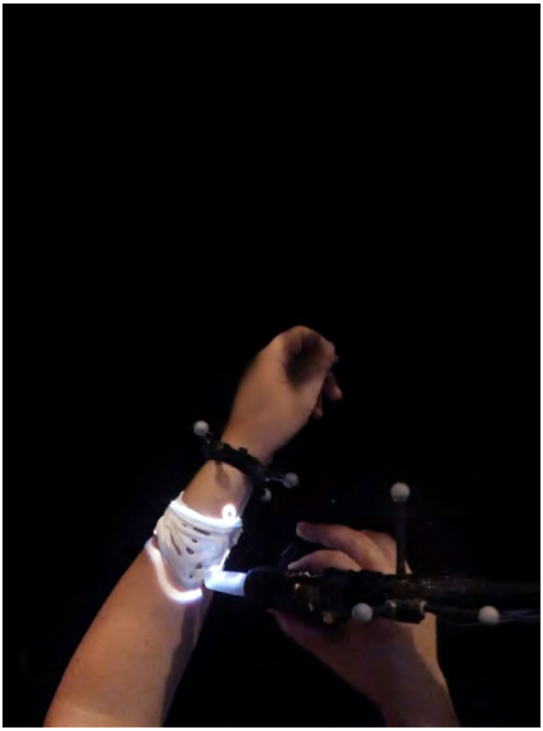

While *ExoSkin* does not produce product-ready artifacts, the process of developing *ExoSkin* was quite useful for eliciting many of the technical challenges unique to coordinating fabrication machines in close proximity to people. As such, this chapter also presents a set of design considerations when creating interfaces for designers and fabrication machines to work in tangible, close cooperation with one another. It details design factors for digital on-body printing, including material choice, machine configurations, appropriate content, and hybrid workflows. Finally, this system demonstrates that fabrication machines built around human affordances can expand digital techniques into previously analog domains. *ExoSkin* shows how integrating machines with context-aware geometry and interfaces can make the line between *tool* and *collaborator* more ambiguous.

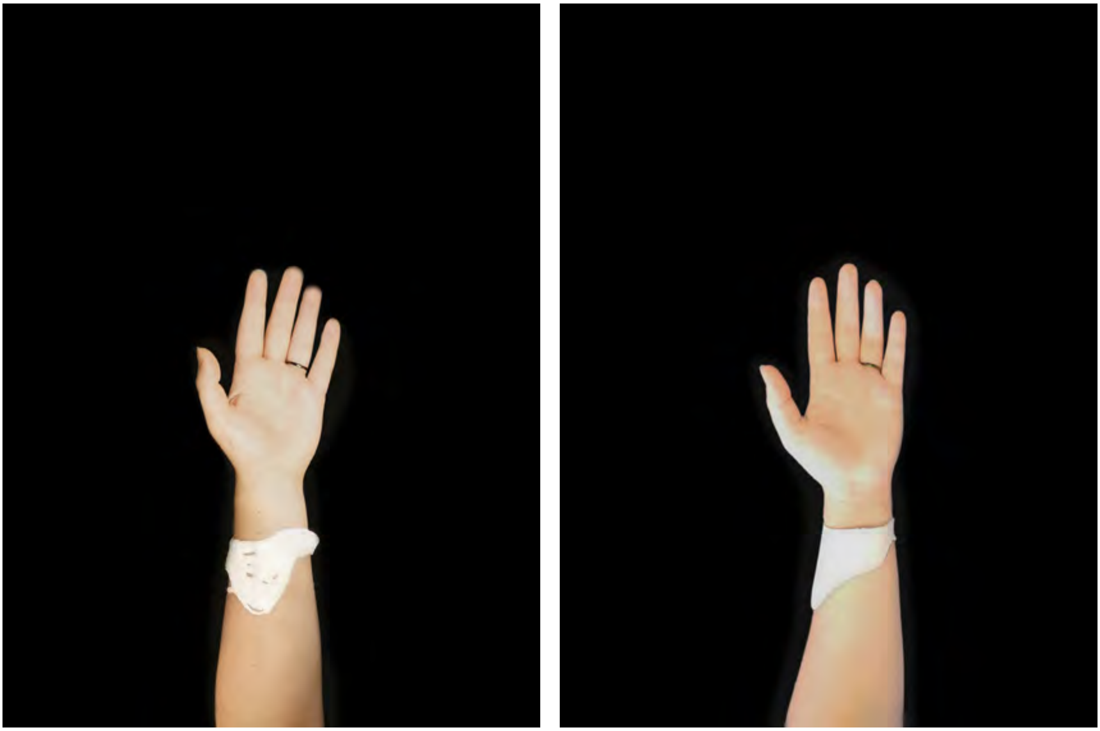

FIGURE 1
ExoSkin is a hybrid fabrication system for designing and printing digital artifacts directly on the body.

On-Body Fabrication

Art and design communities have a strong legacy of crafting wearable artifacts directly on the body or on proxies for the body. Fashion designers and tailors create bespoke garments on models or mannequins; special effects artists craft prosthetics and props on actors or lifecasts; tattoo artists inscribe their graphic designs onto their client's skin. In many of these domains, the techniques for on-body design and fabrication are purely analog. In each of these scenarios, the artifact is customized and hand-crafted on an individual's body by highly trained fabricators. Moreover, whether it is spray-on clothing, self-assembling body suits, or robotic reconstruction of bodies, we see aspirations towards machines making things on the body referenced across pop culture and science-fiction. Despite digital on-body fabrication seeming like a topic of science fiction, this research anticipates a future where wearable artifacts – such as clothing, jewelry, and medical braces – are fabricated in real-time directly on the body through a human-machine collaboration.

However, there are several human, machine, and material challenges unique to on-body fabrication that make it difficult to realize this vision. With traditional Computer-Numeric Control (CNC) processes, material is added or removed on a flat, stabilized build platform. By contrast, when fabricating on the body, the build platform – the body – is a highly deformable, highly curved surface in constant motion. Moreover, materials for on-body fabrication, such as silicones, plasters, clays, or textiles, are dynamic, and are actively transformed by gravity, temperature, and the environment. Last, for safety reasons, traditional CNC processes are not designed for close-quarter human interaction. These requirements cannot be satisfied in on-body fabrication, where the human can be both the operator and the canvas for fabrication.

These challenges have impeded the development of digital workflows specific to on-body fabrication. For example, 3D printers can produce wearable objects, however this fabrication process is very inefficient: the form of objects that wrap the arms, legs, or shoulders tend to have high volumes of space, but low material densities [Gannon 2015]. This high volume-to-density ratio is particularly inefficient in material and time for 3D printing, where a form is sequentially built up layer by layer. By contrast, a printing process specifically designed for on-body fabrication could integrate the body as a three-dimensional support structure, which could reduce time and material wasted in fabrication. Moreover, printing on the body can provide a more direct and engaging user experience, which to date, had not been explored. Finally, a great impediment to on-body fabrication is the amount of disparate knowledge that needs to come together to build a cohesive system.

A core challenge for developing on-body fabrication systems is the diverse range of knowledge necessary for a machine, a material, and a person to work in such close coordination. For example, *ExoSkin* leverages tools and techniques from Computer-Aided Design, computer vision, material science, and hardware design to enable digital models to be fabricated directly on the body.

Design Considerations

This section discusses the innate challenges of developing on-body fabrication systems, and presents a set of considerations required for designing additional on-body fabrication systems.

Fabrication Methods

As described earlier, many crafts exist for manually creating on-body artifacts. With ExoSkin, we wish to adapt these practices by leveraging to benefits of digital tools, such as providing guidance and feedback, and supporting importing and exporting operations.

Computer-controlled fabrication methods insinuate a specific location for the body in relation to the fabrication machine. For example, for 3D printing, the nozzle of the extrusion machine should remain perpendicular to its surface of the body during printing. This spatial limitation prevents traditional 3-axis 3D printers to be used for on-body fabrication. The complex curvature of the body requires a minimum of 4 axes for the extruder to be normal to the printing surface. Machines with 5 or more degrees-of-freedom, such as hand-held extruders or robotic arms, have the ideal flexibility for printing on the body. If hand-held devices are used, it may be hard to reach certain body regions, and there may be challenges in orienting the device to perform the fabrication. These constraints could be remedied if a third-party is performing the fabrication on a subject's body.

Safety

In every case where a fabrication machine comes in close contact with the human body, safety should be a primary concern. The risk of entanglement should be minimized by keeping moving machine parts fully enclosed and away from the body. Pinch points must be avoided by positioning the body in free space and not on a rigid platform. Most importantly, irritations and burns need to be prevented by using skin-safe materials.

Materials

Finding appropriate materials for on-body printing was one of the steeper challenges for this research. There are a number of constraints and limitations that impact the choice of materials for on-body fabrication. As described above, the material must be skin-safe, including its temperature at the time of extrusion. Second, it must be easy to remove and clean, unless the print is meant to be permanent. Third, the material must be resilient to movements and deformation on the body. Lastly, it must be a workable medium for a given fabrication process (e.g. extrudable).

These safety requirements immediately rule out materials that change states using heat. For example, polymers that liquefy with heat, such as the thermo-softening plastics used in Fused Deposition Modeling (FDM) 3D printers, or materials that cure with heat, such as the thermosetting plastics, concrete, resins, clays, and plaster often used in freeform printing, are not appropriate for on-body printing. Photopolymers may be applicable, however they were avoided in this research due to concerns of prolonged UV exposure to the skin.

The second consideration for skin-safe material is that it can be removed from the body. Latex and silicones are commonly used to make prosthetics, masks, and molds for the body, however they often require adhesives to remain stuck to the skin and must be peeled off the body when done. Water-soluble clays and pastes, by contrast, will stick to skin during fabrication, and can be simply washed off after use. Excessive body sweat or humidity, however, may impede drying and can deteriorate the printed material over time.

Other considerations are the specific material properties in relation to a chosen fabrication process. In extrusion, for example, the viscosity of a material plays a critical role. If the material is not viscous enough, it will take a long time to set, and consequently slide off the body during printing. But if the material is too viscous, it will be extremely difficult and slow to extrude. Moreover, attributes like drying rate, drying time, and layer adhesion all influence the performance of the material during extrusion, as well as the finish quality of the fully cured material.

Hybrid Fabrication Workflows

Integrating analog and digital craft into a hybrid fabrication workflow can be strategically challenging. As discussed in Chapter 2, there are a spectrum of possible influences in which the digital tools can control the process, from active to neutral to passive. If the workflow becomes too digitally oriented (active), the benefit of the human agency is minimized. Likewise, if the advantages of digital fabrication processes are not leveraged (passive), the existing analog methods may be limited. In designing hybrid workflows for on-body fabrication it is important to consider where in this spectrum the system should fall.

DESIGN

An important aspect of any fabrication workflow is the design phase. Existing research illustrates techniques for designing digital models on and around the body [Friedman 2014; Wibowo 2012]. Developing additional digital input modes based on tools currently used in on-body fabrication may be a more contextual approach to hybridizing these analog crafts. For example, a tailor could digitally design a garment directly on a customer's body, rather than manually measuring the body dimensions and subsequently producing the design.

Adapting body-based input as a digital process also enables us to augment several existing analog design methods. For example, positioning, scaling, copying, or reorienting a 2D pattern on the body can be a time consuming analog process. However, these geometric transformations are trivial in digital design. Moreover, in many analog on-body processes, such as creating garments or prosthetics, the design and fabrication stages happen simultaneously. The ability to visualize a digital design on the body prior to fabricating can enhance the design-to-production workflow, and enable more rapid design iteration before committing to fabrication.

Finally, one limitation to existing analog on-body fabrication workflows is that the designer and the fabricator are often the same person. This constrains one expert craftsperson to be working on only one subject at a time. However, decoupling design from fabrication in on-body workflows can better leverage the unique affordances of digital processes. Digital tools to support digital workflows have the potential to make on-body designs more easily duplicated, shared, archived, fabricated remotely, or adapted to many different bodies.

FABRICATION

In terms of the actual fabrication process, a fully digital fabrication process could be preferred to ensure the highest level of precision in the final fabricated model. However, there are certain practices that may be best suited as an analog technique within hybrid workflows. For example, the person who will wear the artifact may desire more agency and control over the final outcome. In this case, keeping human input integrated into the design process may increase overall satisfaction, since people have complex sensitivities and preferences for what is worn on their bodies.

Additionally, a fully digitized fabrication process may not be appropriate for hybrid workflows in certain scenarios. For example, fabricating near sensitive or injured areas of the body require a level of delicate and dexterous control that goes far beyond the sensing and actuation capabilities of current fabrication machines, such as robotic arms. In these scenarios, integrating a hand-held or assistive fabrication device into the hybrid workflow may be most effective.

FIGURE 2

ExoSkin's custom fabrication machine is specifically designed for on-body printing. The machine consists of a motorized clay extruder, a hose, and a hand-held extrusion nozzle.

ExoSkin

These design considerations for on-body fabrication were further explored by developing *ExoSkin*, a hybrid fabrication system for printing digital designs directly on the body. This system augments a custom, hand-held fabrication machine so a designer can directly deposit material onto interactive 3D models projected onto their arm. To begin designing with *ExoSkin*, a person first uses the handheld fabrication machine as a stylus, to provide design input directly on the body. The resulting geometry is projected in real time directly on the body, and will stay digitally attached to the designer as they move and rotate their arm. Once satisfied with a design they can switch to an output mode and fabricate the design by extruding a single layer of material. A projected guidance system helps facilitate the fabrication process by visualizing animate toolpaths on the designer's moving body.

Implementation

Hardware Configuration

The hardware configuration of *ExoSkin* features a custom hand-held extruder that is augmented with a motion capture and a projection system. This hardware system enables a designer to digitally fabricate a single layer of material directly on their body.

ON-BODY PRINTER

The main hardware component of *ExoSkin* is its custom fabrication machine designed specifically for on-body printing. It consists of three parts: a motorized clay extruder, a hose, and a hand-held extrusion nozzle (Figure 2). The clay extruder is a Potterbot linear RAM extruder, which has a 2-liter material reservoir. We add an 18" high-pressure polyester reinforced PVC tubing to transfer material from the extruder reservoir to the hand-held effector. *ExoSkin* was implemented using the Java programming language at 60 FPS. It uses the following open-source libraries: Processing for graphics, OpenCV for projection mapping, toxiclibs for computational geometry, and oscP5 for streaming data over

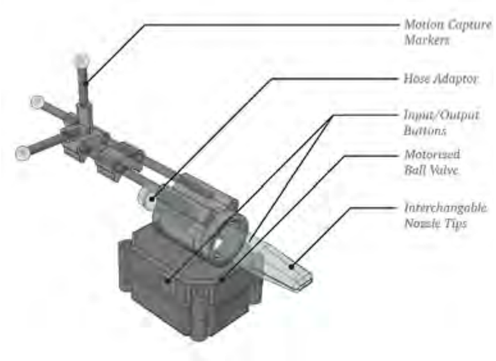

FIGURE 3

Diagram of the hand held extrusion nozzle.

FIGURE 4

The extruder's nozzle tips can be exchanged to modify the flow rate and profile of extruded material.

Open Sound Control (OSC). Rhinoceros 3D and Grasshopper 3D are used to illustrate how *ExoSkin* can send and receive geometry from external CAD software.

The system uses extrusion as its hybrid fabrication method. The decision to develop a hand-held extruder instead of adapting an industrial robot for on-body printing was made over concerns of safety and user engagement. *ExoSkin* is the world's first on-body printing system, and it explores largely uncharted territory in CAD/CAM. Consequentially, developing the additional safety measures needed for a robotic arm to safely touch a human body were beyond the scope of this research. Although a hand-held extruder will be less precise than a robotic arm, users can gain an increased sense of control and agency throughout the design and fabrication process [Devendorf 2015].

END-EFFECTOR

The hand-held portion of the extruder consists of a hose adaptor, motorized ball valve, a set of interchangeable nozzle tips, two input/output buttons, and motion capture markers (Figure 3). The hose adaptor and nozzle tip screw into the motorized ball valve. The motorized ball valve has a slow 3-second phase cycle, which means it takes 3 seconds to fully open or close. However, it can operate with high viscosity material, unlike quicker, but weaker solenoid valves. Prior to fabricating, we pre-pressurize the extruder to push our clay paste from the material reservoir, through the hose, and to the hand-held effector. An output button under the thumb triggers the ball valve to open or close for extruding material. An input button, positioned on the ball valve under the index finger, switches to an input mode which is described later.

The profile of the extruded material can be changed throughout a fabrication session by exchanging nozzles tips on the extruder (Figure 4). Swapping a small diameter for a large diameter nozzle will help rapidly increase the volume of material extruded. Swapping a low perimeter tip for a high perimeter tip will improve drying times for our air-dry clay, since increasing the surface area-to-volume ratio exposes the extruded section to more air.

Many factors contribute to the flow rate of the material from the extruder nozzle: the viscosity of the material, the shape of the nozzle profile, and the length of tube between the material reservoir and the hand-held effector. In our system, the flow rate would vary from approximately one to four inches per second, mostly depending on the viscosity of the material and the size of the nozzle being used.

Guidance System

ExoSkin uses a fabrication guidance system for tracking and visualizing toolpaths data directly on the body. The guidance system is comprised of a motion capture tracking system, projection-mapped visual feedback, and user input.

MOTION TRACKING

ExoSkin uses a six camera OptiTrack motion capture system mounted above a 3'x 4'x 3' tracking area. Motion capture cameras are mounted from above and below in order to track the full rotation of the arm. Passive markers attached to the wrist of the user track the position and orientation for the arm in world space. Markers attached to the end of the extruder track the position and orientation of the nozzle tip.

Motion capture is handled by OptiTrack's Arena (version 1.7.3). It tracks two pre-defined rigid bodies: the wrist worn marker set and the marker set on the extrusion tool. The position and orientation of these markers are streamed over Open Sound Control (OSC) to our software that controls body-based input, output, and projection mapping.

VISUAL FEEDBACK

Generic toolpaths used by traditional CNC machines are sometimes visualized in software as simple lines that map how a tool head will move across a volume of space. However, the body is a more complex canvas for fabrication, and brings a number of complications for the generic toolpath. To begin, the user is printing with a hand-held extruder, which is inherently less accurate than a CNC controlled extruder. Moreover, the extruder must move relative to the arm; not a volume of space. As a result, this highly curved, constantly shifting surface will have parts of toolpaths that go around the body, and can't be seen by the user. Our system projects custom toolpath visualizations directly on the body to mitigate these challenges unique to on-body fabrication.

A DLP projector is mounted above the tracking area to provide on-body visual feedback. To accurately project onto the body, we first calibrate the projector to the motion capture system using a simple one-time routine that correlates projected 2D points with tracked 3D world points.

FIGURE 5

Motion capture markers placed on the hand-held extruder and the designer's arm enables helps ExoSkin accurately visualize user input on their moving body.

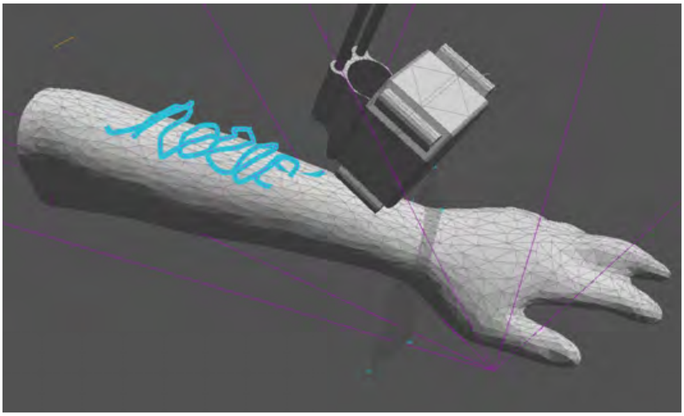

FIGURE 6

The CAD backend in ExoSkin synchronizes the designer's body, the fabrication machine, and the physical work space into a unfied virtual environment.

User Input

In addition to fabrication, the hand-held extruder also serves as a digital input device. The user enters an input mode by pressing the input button on the extruder. In this input mode, the user can draw digital content on their physical arm by using the tip of the extruder as an on-body stylus (Figure 5).

A CAD backend records the world coordinates of the extrusion nozzle's tip. These user-recorded points are down sampled and smoothed from the initial motion capture data. The motion capture system streams coordinates at 100 fps at sub-millimeter precision. The filtered coordinates are then attached to the virtual representation of the arm. This lets the user freely move and rotate their physical arm, while keeping the user input geometry fixed to its original location relative to the reoriented arm.

CAD Backend

ExoSkin has a CAD backend that is not seen by the user. It holds the pre-scanned mesh of the body part intended for on-body fabrication, a mesh of the extrusion tool, and the 3D model of the current design. The meshes are dynamically rigged to the rigid body data streaming from the Arena software. The CAD backend aligns the arm's mesh to the incoming tracking data by translating the mesh from a known offset to the incoming marker coordinate, and then reorients the medial axis of the mesh to the normal of the wrist-marker plane. Similarly, the mesh and tip of the extrusion tool is transformed from a known offset to the position and orientation streaming from the tool-marker set. Synchronizing a virtual representation of the arm and extrusion tool with the physical arm and extrusion tool enables the CAD backend to record user input in coordinates that are relative to the moving body (Figure 6).

FIGURE 7

ExoSkin uses a water-based polymer clay as its printing medium.

FIGURE 8

This timelapse shows ExoSkin's extruded air-dry polymer clay after 0 minutes, 2 minutes, and 6 minutes. The clay dries with a smooth surface finish and a foam-like flexibility.

Material Configuration

ExoSkin initially explored three types of clays as a possible printing medium: an oil-based polymer clay, a water-based polymer clay, and a stoneware clay. The decision to focus on clays as a printing medium was guided by the constraints identified in the design considerations section of this chapter. The viscosity of each substance is controlled by adding a thinner to the first oil-based polymer clay, and water to the second polymer clay and stoneware clay. Each material was extrudable, however the oil-based polymer clay would not harden at room temperature. The stoneware clay would harden too quickly, and was prone to cracking and chipping.

In the end, the water-based polymer clay (Jumping Clay) was the most reliable printing medium to work with (Figure 7). It has a few unique affordances that make it an ideal candidate for on-body printing. First, it is an air-dry clay, so it cures from liquid to solid at room temperature. Second, when cured it is a lightweight, semi-flexible foam (Figure 8). This flexibility is an ideal material property for printing on the body, as the surface finish is resistant to cracking as the skin deforms. Lastly, this particular clay is reusable. Even when fully cured, this clay can be harvested and returned to its paste-like state by submerging it in water. This last material attribute is unique and compelling affordance, as the majority of materials currently used in 3D printing are either one-time use only, or have a complex recycling process.

FIGURE 9

The user holds down the output button and extrudes over the rendered design.

User Interaction and Workflow

ExoSkin strategically implements a hybrid fabrication workflow that enables digital designs to be printed directly on the body. Its hybrid workflow builds upon existing analog techniques in on-body fabrication. The analog design process for on-body fabrication happens *in-situ* -- within the local context of the body. This enables an analog fabricator to dynamically adapt a design in response to the surface, the person, or other conditions that would otherwise impact the final outcome. *ExoSkin* adapts this dynamic, in-situ design process as a digital workflow. The designs which are created are meant to demonstrate the capabilities of the system, and do not necessarily represent product-ready models.

DESIGN

A user begins by digitally drawing on the skin by pressing and holding the input button on the extruder. As they move the extruder over the arm in real-world space, the sketched lines are projected as toolpaths directly onto the body (Figure 9). Together, the tracking system and projection mapping keeps a persistent rendering of the artwork relative to the moving arm. This means that when the user moves or rotates their arm, only the correct, visible portions of the design are rendered. The system down samples motion capture data to smooth user input to a minimum distance of 3mm between points. Additionally, *ExoSkin* automatically snaps user input points to the closest points on the virtual surface. These filtered surface points are then interpolated into a smooth 3-degree spline.

FABRICATE

Once a sketch is complete the user can transition to the output mode of the extruder. *ExoSkin* illustrates on-body fabrication by focusing on the arm, since this body part is highly mobile, has a complex curved form, and tires quickly when unsupported. However, the techniques and ideas explored on the arm are largely applicable to other parts of the body.

ExoSkin uses a neutral hybrid fabrication process: digital guidance is provided to the user, but there is no actuation

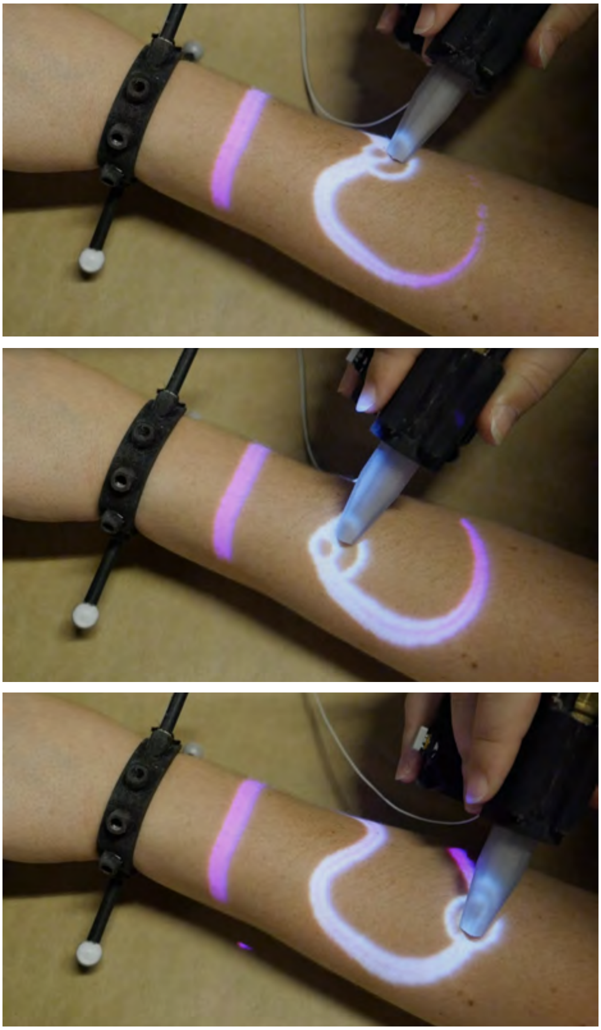

FIGURE 10

ExoSkin's projected guidance system provides continuous visual feedback for the tool relative to the body. Here, it renders a dynamic toolpath that indicated proximity and directionality to the user..

or force feedback to control or constrain the geometry as it is fabricated. A user holds down the output button on the nozzle to begin the fabrication process, which opens the ball valve to begin extruding material. This fabrication process is facilitated by a projection guidance system which visualize the required tool paths: the user slowly traces the projected toolpaths rendered on their body to print a design.

ExoSkin's projected toolpaths provide multiple layers of visual information and feedback to assist the user during the fabrication process. These interactive, animate toolpaths help improve accuracy when fabricating with the hand-held device. For example, *ExoSkin* provides continuous, visually augmented feedback about the nozzle location in relation to the toolpath: as the user brings the extruder near a toolpath, a halo around the tip of the extruder and the closest point on the toolpath is projected (Figure 10). The halo illustrates the disparity between locations and visually prompts the user to adapt their tool position.

Additionally, projected toolpaths in *ExoSkin* provide subtle feedback how to reorient the arm for fabricating on non-visible parts of the body. For example, the line thickness and color gradient of the toolpath are dynamically animated to grow thicker and brighter directly under the extruder. However, as the user moves the extruder tip from one side of the body to the other, the thickness and color of the toolpath thins and dims, as if it is wrapping around to the backside of the surface. This discreet tapering technique enables *ExoSkin*'s 2D projection plane to guide extruder paths around a highly three-dimensional surface.

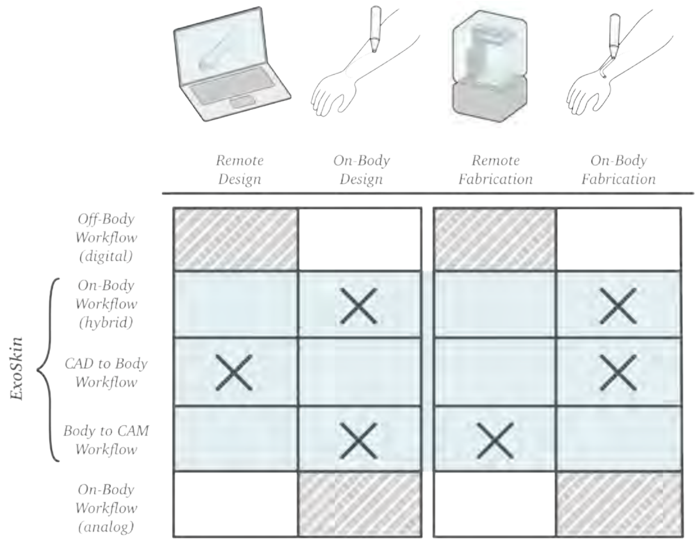

FIGURE 11

ExoSkin enables three new workflows to better support analog-digital ways of making things for the body.

FIGURE 12

This sketch beautification demo shows ExoSkin taking imprecise input to create precise geometic output.

Designs generated through *ExoSkin* are simple and abstract: the emphasis of this system is less on generating product-ready models, and more on demonstrating the feasibility of on-body fabrication. Moreover, *ExoSkin* supports multiple digital-analog workflows for designing and fabricating on the body (Figure 11). Since its guidance system fluidly tracks physical interactions between designer and machine, *ExoSkin* can map movements from the designer's body and fabrication machine into a unified virtual environment. For example, this persistent tracking enables the system to simultaneously construct a digital model of a wearable as the designer sketches the input on their body. The designer then has the option of continuing to the fabrication process on the body or remotely, using traditional 3D printing. Similarly, a designer can carefully draft a design in a traditional CAD program, and then use *ExoSkin* projected guidance system to fabricate it on the body.

Another unique advantage of this hybrid fabrication process is that users can physically manipulate the digital design, if they do not get a desired result once the material is extruded. The flexibility built into *ExoSkin* allows minor errors to be corrected just by nudging the material to better match a desired toolpath. Larger errors can be corrected by removing portions of the fabricated path and re-printing it.

Sketch Beautification

ExoSkin also implements simple sketch beautification [Igarishi 1997] to transform imprecise user input into precise geometric objects. In Figure 12, the user quickly sketches a circular shape onto their body using the input mode of the extruder. The system then recognizes the sketch as a circle, creates the idealized geometry on the virtual arm, and then projects the 3D dimensional perfect circle as a toolpath, and the user can press the output button to extrude material while tracing the toolpath. Likewise, a triangle is sketched on the body, processed into an idealized polygon, and then projected back as a precise shape. Again, the user presses the output button on the extruder and traces the toolpath to fabricate the shape.

FIGURE 13

Designs in ExoSkin can be generated in commercial CAD software and then projected onto the body for fabrication.

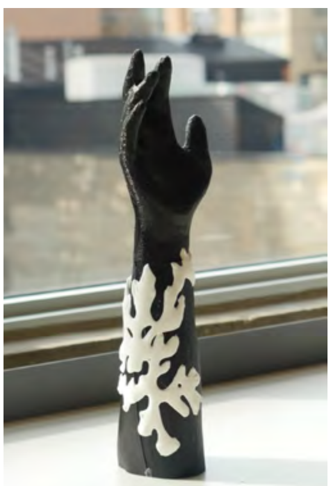

FIGURE 14

Decoupling input and output means that the final design can be fabricated locally on the body or remotely across many bodies.

A Framework of On-Body Fabrication

ExoSkin provides several new, alternative pathways for analog and digital on-body fabrication. Since this hybrid fabrication system uses its on-body printer an input and output device, and since it always knows where it and the body are in space, *ExoSkin* can import existing geometry from conventional CAD software or export constructed geometry to conventional 3D printers. This interoperability between digital-analog design and fabrication enables *ExoSkin* to augment previously unsupported analog crafts with human-centered digital tools. Figure 11 contrasts the input and output workflow paradigms facilitated with ExoSkin, to existing workflows for the design and fabrication of wearable objects. The flexibility *ExoSkin* provides to mix input and output methods lays the groundwork for a number of opportunities worthy of future exploration.

Importing Geometry

For on alternative workflow, *ExoSkin* connects to conventional 3D modeling software to import pre-made geometry. This lets a user create precise digital designs for the body, which are then sent to our system for on-body projection and fabrication. In Figure 13, a user designs an organic 2D pattern in Rhinoceros3D, a commercially available CAD software (Figure 13a). I developed a script running in the CAD program to map the 2D pattern to the 3D mesh stored in our system, and then send the 3D geometry to ExoSkin via OSC. *ExoSkin* stores the geometry as toolpaths attached to the virtual arm. This keeps the CAD-generated geometry attached to the user's arm as they move and rotate around the workstation. These CAD-generated geometries are projected onto the body, and are ready for printing (Figure 13b).

ExoSkin's pipeline to existing 3D modeling software enables users to quickly test out a design directly on the body. Users can easily translate, rotate, copy, or scale a design in the CAD program, which then updates the geometry projected onto the body in real-time. The floral pattern armlet in Figure 13b demonstrates *ExoSkin*'s capability of fabricating highly intricate toolpaths. Furthermore, Figure 14 demonstrates how *ExoSkin* facilitates digital designs being physically fabricated on other people's bodies.

FIGURE 15

ExoSkin *supports several analog-digital workflows for fabricating on the body.*

Exporting Geometry

ExoSkin's connection to external CAD software brings an additional benefit to the hybrid fabrication process: the ability to export an artifact designed or fabricated on the body as a ready-to-print 3D model (Figure 15). In Figure 15a, the extruder is used as an input device to draw the design of a bracelet directly on the arm. ExoSkin stores the user input as a set of toolpaths, which it also sends to a connected CAD program. The user then prints the bracelet on the body by tracing the projected toolpaths, then infilling the defined region (Figure 15b). Simultaneously, a script running in the CAD program generates a thickened surface from the user's input geometry, and exports a valid mesh for conventional 3D printing (Figure 15d). This allows users to obtain a quick physical prototype printed directly on the body, and then print a more robust model on an external 3D printer with rigid 3D printing materials.

Initial Observation Session

ExoSkin is an exploratory system designed to investigate the concept of on-body fabrication. As such, this observation session is not intended to evaluate the system or conduct formal caparisons to other fabrication techniques. However, it is still useful to get initial user feedback on the system and the concepts which it represents.

Four users participated in a single workshop for this observation session: two male and two female, aged 20–30. Participants were internal to our organization but were not members of our research group, and had no prior knowledge or exposure to the ExoSkin system.

All four participants had engineering or computer science backgrounds, but had very little previous experience with digital design and fabrication methods. One participant had basic knowledge of 3D modeling and only one participant had ever used a 3D printer.

Procedure

During the workshop session, I first explained the system configuration, providing an overview of the fabrication, sensing, and projection technology. The fabrication workflows were then demonstrated to participants, highlighting the main features of the system. Each of the four participants were then given an opportunity to use the system, and were prompted to design and fabricate a simple model directly on their arm. The entire session took approximately 45 minutes. Below I discuss the main observations and feedback which were collected.

Participant Feedback

In general, participants reflected positively on our system and were enthusiastic about on-body fabrication. When initially extruding onto the skin, each participant made a comment on the printing material's texture. They were surprised by its stickiness and cool temperature, with one participant exclaiming "*It tickles!*" (*p*2).

Participants did see the value in printing with a washable, reusable material, however there was also a desire to preserve a print so that it can be worn multiple times:

> *"I would like to be able to take it off my body without destroying my beautiful design."* (p1)

When prompted for their thoughts on the extrusion tool, three participants noted a potential preference for a robotic instead of hand-held tool. They felt a robotic system would increase the speed and accuracy of the fabrication process and would take less effort on their part. There was also a sense that a machine would do a better job than they themselves could:

> *"I don't trust myself to make it right… I'd rather trust a machine."* (p3)

P4 enjoyed the hand-held device, and also noted the desire to use it both on his body and on others:

> *"I'd use it to give someone a tattoo!"* (p4)

Participants were also asked if they would entrust a machine to fabricate on sensitive areas of their body, such as the face or back. Three participants gave a definitive 'No', however one participant said they would feel comfortable as long as there was a human overseer:

> *"I wouldn't mind a machine printing on my face … as long as a person would come check up on me every once in a while. You know … to check if I'm still alive."* (p2)

At the end of the workshop, participants were asked to think about the kinds of thing they would print on their bodies. Clothing and accessories were immediately mentioned, and one participant wanted to print custom electronics:

> *"I'd print a game controller right on my arm!"* (p3)

Discussion

There are many analog craft processes that can be adapted for hybrid workflows in on-body fabrication. *ExoSkin* illustrates an optimism for future explorations into on-body fabrication. Many diverse and innovative interaction techniques will be needed to circumvent the unique challenge of crafting digital designs directly on the body.

Domain Specific Applications

This chapter focuses on the interaction implications and hardware configurations for printing on the body. Although *ExoSkin* provides some examples of the artifacts which could be generated with such a system, further work could investigate specific applications for fabricating functional objects on the body.

Body-centric design industries that rely on one-off, handcrafted designs are currently limited in their ability to rapidly create, iterate, and share a given design. Augmenting these high-skill analog craft practices with digital techniques brings an opportunity to streamline design-to-production workflows for on-body fabrication. There are four primary domains that would greatly benefit from on-body digital fabrication: skin-centric industries (such as cosmetics), fashion and wearables industries, film and special effects industries, and the medical device industry. Such domains may specifically benefit from on-body design, on-body fabrication, or both.

In applications that modify the skin, such as cosmetics, digital drawing or brushing instructions that are projected onto the body can help non-experts learn expert techniques. In the movie and special effects industries, body-worn props or prosthetics can be precisely designed in a CAD environment, then digitally fabricated directly onto an actor. For fashion and wearable devices, on-body design could give both the artist and the model agency in a customized design process, while direct-to-body fabrication enables bespoke designs to be customized to many bodies, and inherently ensures the design will fit the wearer. Lastly, for medical applications, the design of prosthetic sockets or silicone dressings, for example, can be created by an expert then sent to remote locations for technicians to fabricate directly on a patient.

Body Parts

ExoSkin focuses on a single body part: the arm. The arm as a representative body part, as it is highly mobile, if has a complex curved form, and it tires quickly when unsupported. However, the earlier discussion of possible content domains indicates that many areas of the body may be appropriate for on-body fabrication. However, each specific location brings unique challenges and considerations that impact the choice of machine configuration. This can relate to physically sensitive areas of the body – such as the face, head, neck, and spine – socially sensitive areas – such as the chest and nether regions – or highly mobile areas – such as the arms, hands, legs, and feet.

Printing on physically sensitive areas requires continuous, nuanced feedback on how the fabrication device is touching the body. In these scenarios, hand-held devices may be most reliable and appropriate as fabrication devices. Socially sensitive parts of the body, may be less desirable for direct on-body fabrication or may be preferred to be operated by the person being printed on instead of a third party. For highly mobile areas of the body, the fabrication system cannot assume that the body part will remain still for long periods of time. Therefore, the configuration of the system must be designed to adapt to continuous changes in position and orientation of the fabrication surface.

Limitations and Future Work

ExoSkin deliberately implements a hand-held extruder as the fabrication device for this on-body fabrication system. This decision was guided by the relatively quick development time for a handheld device, as well as the low risk of injury for the end user. However, what was gained in agile deployment and increased safety, was lost in precision and accuracy when compared to existing multi-axis CNC machines. Future work could focus on developing compliant, multi-axis systems that strike a more even balance between precision, accuracy, and safety for on-body fabrication. For example, worker-friendly robotic arms, such as collaborative or soft robotic arms, would be particularly interesting to explore.

ExoSkin uses a motion capture system as a low-latency, high accuracy, and highly flexible solution to track bodies in space. However, there are other spatial tracking technologies that could be useful for on-body fabrication systems. Markerless tracking, in particular, is most desirable for any sort of deployment. Early development for *ExoSkin* experimented with both Kinect and Leap Motion sensor, but found that the accuracy of these camera systems was not reliable enough for on-body fabrication.

ExoSkin uses a single projector to visualize fabrication instructions on the body. Although the throw of the projector adequately covered the volume of our work area, shadows cast by the extrusion tool could hide portions of projected toolpaths from the user's view. Future fabrication systems could mitigate this problem by using alternative visualization configurations. For example, switching to multiple projectors or using augmented reality devices, such as translucent screens or head-mounted displays, would eliminate the shadows cast by physical objects in the workstation.

ExoSkin examines the implications of direct on-body printing using a single, skin-safe material. However, there are many more materials to be developed and explored. We are particularly excited for composites that layer skin-safe and non-skin-safe materials together. For example, a skin-safe paste could be printed as an insulating layer against other heat-transferring materials, like thermo-softening or thermosetting plastics. Moreover, edible materials such as frostings, pastes, or foams may also be applicable for on-body fabrication.

One notable limitation of this implementation is that it supports extrusion of only a single layer of material. While *ExoSkin* supports complex toolpaths that are curved and three-dimensional geometry, it does not provide traditional multi-layer three-dimensional fabrication. To support this, the set time of the material would need to be accounted for. Furthermore, or toolpath generation algorithm would need to be advanced to support multi-later digitization.

On-body fabrication could also be extended to support printing on areas of the body under high amounts of deformation or stress, such as joints, hands, and feet. Moreover, fabricating electronics directly on the body could be explored by combining skin-safe materials and conductive pastes or inks. The machine processes for on-body fabrication also warrant further investigation. *ExoSkin* uses a custom material in a CNC extrusion device. However, existing everyday materials could prove useful if the appropriate fabrication process were developed. For example, threads and textiles could be used in fabrication machines that weave or drape directly on the body, and medical tapes or gauzes could be used in machines that wrap bandages, braces, casts, or splints around patients.

Fabricating with a custom-made material also limits quality control from batch to batch. Despite our best efforts, the water content of each hand-mixed batch of polymer clay would vary slightly. As a result, the behavior of the material, its viscosity and drying time, would differ each time the material reservoir was reloaded. Future implementations could integrate an air-assist onto the extrusion tool to actively dry the clay when too wet.

Finally, the preliminary observation session with *ExoSkin* provided some interesting insights. For example, the issue of reusability came up – how can models printed directly on the body be preserved and re-worn? However, more thorough evaluations about the implications of on-body fabrication could be conducted. Furthermore, the topic of human agency revealed subjectivity and trade-offs: while some identified the manual extrusion tool as potentially imprecise, others valued the ability it gave them to control the fabrication process. Follow-up studies could explore these and other related topics.

Conclusion

This chapter demonstrates how existing body-centric analog crafts can be augmented with hybrid fabrication workflows to enable digital designs to be crafted directly on the body. An initial observation session about on-body fabrication that was conducted with *ExoSkin* indicated that this new paradigm for interactive 3D modeling involves curiosity and intrigue, motivating further explorations and implementations. Moreover, the design considerations for future on-body fabrication systems outlined in this chapter identify the unique human, machine, and material challenges that these systems will need to solve. Digital on-body fabrication is an admittedly challenging domain. However, the potential impact to augmented these previously analog crafts with digital workflows is an important area for future explorations.

Chapter 6

Mimus

Eliciting Life-like Qualities from Robotic Arms

Reverb, *Tactum*, and *ExoSkin* each contributed critical insights into *why* and *how* to design human-centered interfaces for autonomous fabrication machines. *Reverb* demonstrated that expert knowledge from a fabrication process can be embedded into interactive geometry. It also provided critical insights for integrating autonomous behaviors into a virtual geometry. *Tactum* then brought this intelligent, animate geometry out of the computer to interact in the real, physical world. Its fabrication-aware and context-aware interface was my first experience exploring how virtual geometry could dynamically engage with people moving in space. *ExoSkin* then worked to close the loop between design and fabrication, and connected intelligent, on-body virtual geometry to a custom, hand-held fabrication machine.

ExoSkin revealed the challenges, limitations, and new possibilities of a technology, a person, and a machine working synchronously with one another. This system fleshed out the level of intelligence needed for a fabrication machine to have safe close-quarter interaction with a person. However, as a hand-held device, the fabrication machine I developed was limited to functioning like a tool. However, through this project I saw the potential for a machine that could become an attentive contributing partner. This limitation inspired the next phase of my research, in which I focused more extensively on human-centered interfaces for autonomous robotic arms.

Robotic arms are incredible fabrication machines due to their speed, precision, reliability, strength, and adaptability. Most fabrication machines are built for a single purpose. However, robotic arms can do an infinite number of tasks just by changing its *end-effector* – the tool on the end of a robot. Moreover, these machines act as a bridge between virtual and physical worlds: their minds are in the virtual, but their bodies are in the physical. However, they are not without limitations. Robotic arms are intended to execute short, pre-programmed tasks, and are not well configured for reacting to changing environments or for open-ended control. Moreover, robotic arms – especially industrial robotic arms – can be incredibly dangerous to work with. These machines can operate in the physical world, but they cannot see it. My experiences with *ExoSkin*'s interaction shortcomings motivated me to explore how industrial robots could be reconfigured to see and respond to people in real time.

FIGURE 1
Mimus is a 1,200kg industrial robot that has been transformed into a living, breathing mehcanical creature.

In this chapter, I highlight my early experiments in designing human-machine interfaces for industrial robots. This set of projects reveal insights into how I began thinking and making interactive systems for these arcane machines. My early explorations with robotic arms progress from discovering their innate affordances and limitations, to developing lower-level control software for interactive applications, to prototyping possible interactions and use cases for creative industrial robotics. With this background work documented, I then present *Mimus*: a 1,200kg industrial robot I transformed into a living, breathing mechanical creature. *Mimus* is the culmination of my inquiry into human-centered interfaces for fabrication machines. This work synthesizes the design and interaction techniques from *Reverb*, *Tactum*, and *ExoSkin* into an interactive installation that demonstrates how a well-designed, human-centered interface can foster new, empathic relationships between people and intelligent fabrication machines.

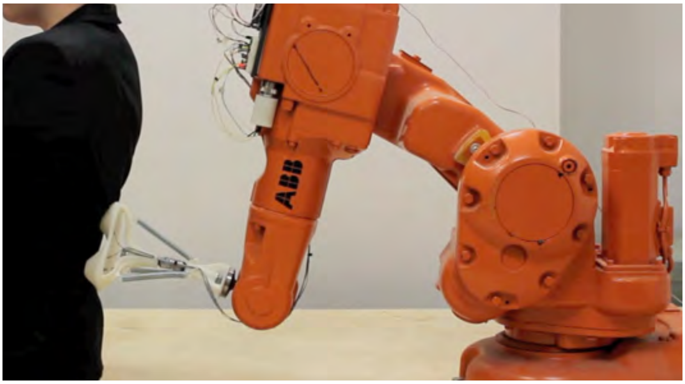

FIGURE 2

RoboMasseuse addresses common taboos in industrial robotics to enable back-controlled interactions between a user and robot.

FIGURE 3

Robo.Op is an open-source hardware and software platform built to open industrial robots to new, more creative applications.

Early Explorations in Robotics

From my onset of working with industrial robots, I have looked beyond how these machines are used today, to explore new ways of leveraging their unique abilities for non-automated purposes. This section documents a series of smaller projects and toolkits I developed, which were formative the development of my sensibilities towards human-robot interaction. These projects provide technical and conceptual insights into how industrial robots can become easier and more intuitive to use. Moreover, they hint at new opportunities for a-typical use cases for these machines in more diverse areas of inquiry.

ROBOMASSEUSE

RoboMasseuse (2012) reconfigures an industrial robot to give safe, sensual back massages to users (Figure 2). Made in collaboration with Zack Jacobson-Weaver, this back-based interface encourages people overcome their anxieties by coming in direct physical contact with industrial robots. This project was important for several reasons. To begin, it was my first experience with an industrial robot. Additionally, it pushed me to speculate beyond the intended purpose these machines today. Finally, and most pragmatically, it prompted me to develop layered strategies to bypass the safety limitations of industrial robots.

ROBO.OP

The code base for *RoboMasseuse* was then formalized into *Robo.Op*, an open hardware / open software platform for hacking industrial robots [Bard 2014]. The toolkit consists of a modular physical prototyping platform, a simplified software interface, and a centralized hub for sharing knowledge, tools, and code (Figure 3). *Robo.Op* translates the obscure machine language of ABB industrial robots (RAPID) into a more modern and accessible arts-engineer toolkit (Processing). This toolkit aims to help non-roboticists engage with industrial robotics, and was presented at the 2014 Open Hardware Summit.

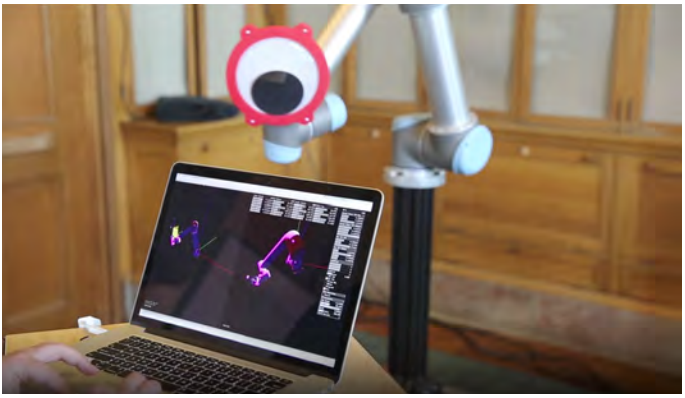

FIGURE 4

ofxRobotArm brings a slew of example projects demonstrating the most common strategies for engaging with collaborative robotic arms.

FIGURE 5

Quipt enables an industrial robot to see and respond to people within the same shared space. The fluid feedback of this interface illustrates the potential for close-quarter interaction with dangerous fabriction machines.

OFXROBOTARM

These ideas were revisited and expanded upon two years later with *ofxRobotArm*, another open source library for creative robotics. Developed in collaboration with Dan Moore, *ofxRobotArm* enables users to control a collaborative robot arm in openFrameworks, a C++ based arts-engineering toolkit. Like *Robo.Op*, the goal of *ofxRobotArm* is to remove the many technical barriers that impede newcomers when first working with 6-axis robot arms. However, this library also included robust examples for the four most common methods of controlling and interacting with a robot arm: direct manipulation, geometry-based manipulation, motion capture-based interactions, and keyframe animation (Figure 4). This project aimed to more specifically encourage others to extend human-robot interaction in new and diverse ways, and was presented at the 2016 Open Hardware Summit Europe.

Quipt

My first series of robotics projects all focus on opening the software and hardware that opens industrial robots for more non-traditional uses. These experiences improving the usability and safety of robotic arms eventually led towards a more ambitious project, *Quipt*. *Quipt* is a gesture-based interface that lets an industrial robot see and respond to people in a shared environment (Figure 5). Quipt connects a motion capture system – which can see, but not move in the world – to an industrial robot – which can move in, but not see the world. *Quipt* augments the robot with the submillimeter precision of a motion capture system, and enables safe, close-quarter interaction between a person and a giant industrial robot. Moreover, its fluid feedback between the persona and robot's movements facilitated a new, more intuitive way of communicating.

Quipt directly preceded and inspired *Mimus*. Like *Mimus*, *Quipt* augments an ABB IRB 6700 industrial robot by giving it *eyes* into its environment. However, its sensing system is quite different. *Quipt* uses wearable markers that are tracked by a Vicon motion capture system, to see how a person is moving through a space. When these markers are worn

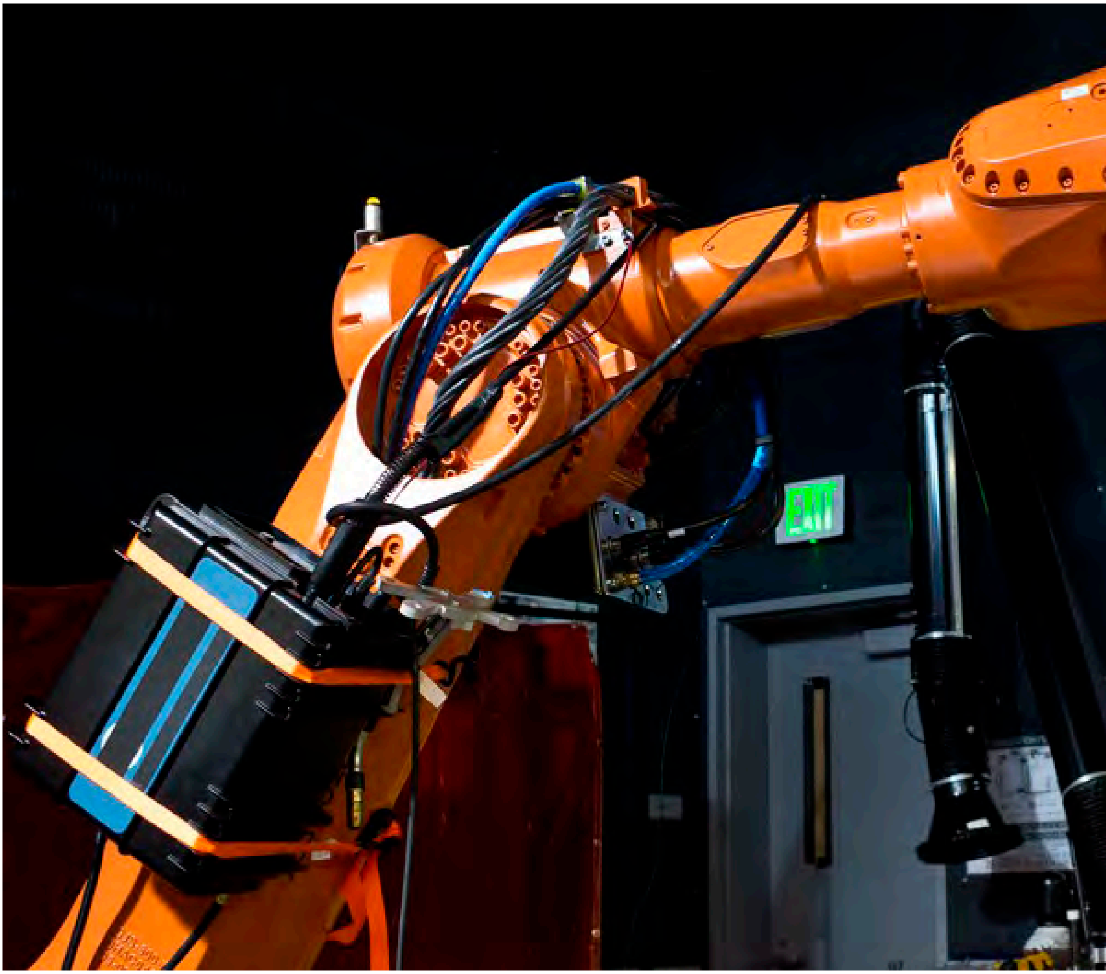

the hand, around the neck, or elsewhere on the body, they enable the robot to safely follow, mirror, or avoid a nearby person.

Quipt enabled the robot to move and track a person with sub-millimeter accuracy. This new level of precision and responsiveness made it possible for a person to safely interact inches away from this giant, dangerous machine. However, *Quipt*'s markered tracking system has its own limitations. To begin, the robot cannot see surfaces or objects that are not attached with markers. While this tracking method is adequate for highly controlled settings, it is not suitable for live, unpredictable environments. Furthermore, markered tracking is limited by the number of

FIGURE 7

objects it can track: if there are too many objects in too small a space, occlusion and disambiguation issues can occur. This limitation prevents the robot from safely tracking a crowd of people, for example.

Quipt directly connected a person's body to the robot's movements (Figure 7). More importantly, it provided a fluid, intuitive communication pipeline for this highly technical piece of machinery. This *literal* human-centered interface created the new ability to rapidly experiment with interactions and experiences for the robot. The intimacy and familiarity I consequentially developed with this robot was further developed, refined, and formalized in my next project, *Mimus*.

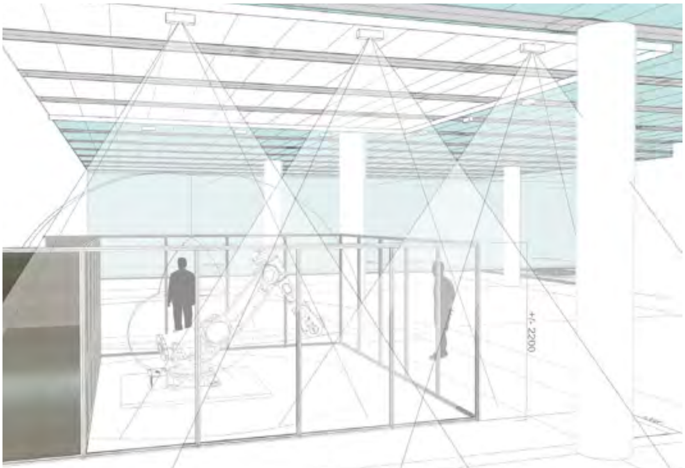

FIGURE 8

Mimus lives behind a clear enclosure, which allows visitors to safely interact inches away from her.

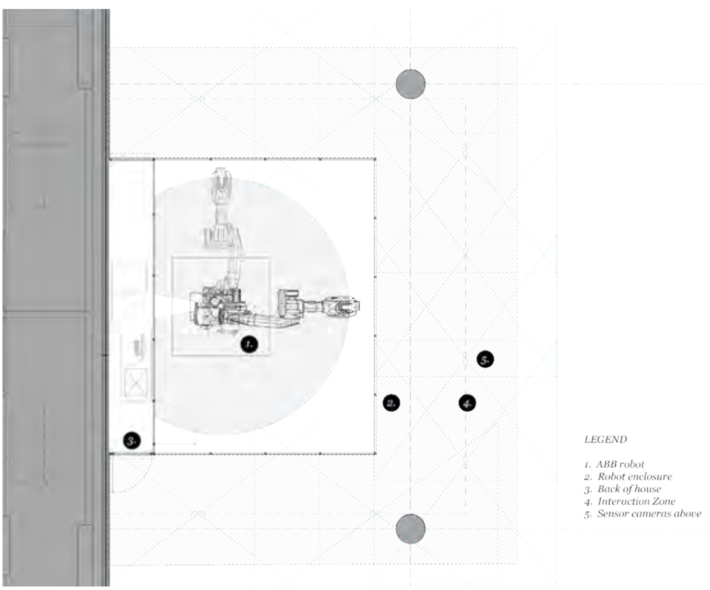

FIGURE 9

The floor plan of the installation encompasses a 9m x 9m x 4.5m space.

Mimus

Mimus is a 1,200kg industrial robot that has been transformed into a living, breathing mechanical creature. More than a tool for performing repetitive tasks, she is able to sense and respond to your presence as you near her enclosure. Unlike in traditional industrial robotics, Mimus has no pre-planned movements: she is programmed with the freedom and curiosity to explore and roam about her enclosure. Mimus has no eyes, however – she uses sensors embedded in the ceiling to see everyone around her simultaneously. If she finds you interesting, Mimus may come in for a closer look and follow you around. But her attention span is limited: if you stay still for too long, she will get bored and seek out someone else to investigate.

Mimus lived at the Design Museum London from November 24th – April 23rd, 2017 as a part of its inaugural exhibition, *Fear And Love: Reactions to a Complex World.*

Implementation

Hardware Configuration

Mimus is comprised of an ABB IRB 6700-200/2.6 industrial robot, an enclosure made of ABB Jokab Quick-Guard machine fencing, and a custom, computer vision-based people tracking system (Figure 8). The installation occupies a physical space of 8m x 8m x 4.5m. *Mimus*'s people tracking system is made of (8) first-generation Microsoft Kinect depth cameras mounted 4.5m above the ground. Each camera is positioned with a 1m overlapping region with its neighbor. This sensor arrangement provides an effective tracking region that is 1.5m wide, running along the entire perimeter of *Mimus*'s enclosure (Figure 9). The tracking system can detect object that are between 500 – 2200 millimeters off the ground. This brings the total footprint of *Mimus* to 11m x 9.5m x 4.5m. The eight ceiling-mounted cameras are connected to a custom-built PC stowed in the back of house area. The PC features three dedicated 4-Port quad bus PCIe USB 3.0 card adapters, an Intel Core i7 6-Core 3.6 GHz processor, an ASUS GeForce GTX 1070 graphics card, 32GB (4 x 8GB) DDR4 RAM, and an ASUS X99-A motherboard.

FIGURE 10

Although Mimus is limited to a birds-eye view of her visitors, the eight ceiling-mounted depth cameras can accurately locate the head height of visitors.

Software Configuration

Mimus uses three layers of custom built software to transform an ABB IRB 6700 industrial robot into an attentive mechanical creature: a people tracking, interaction, and robot control system. The people tracking system handles all the data streaming from the eight depth sensors embedded in the ceiling to the custom PC. The interaction interprets data from the people tracking software to influence the autonomous behaviors of the robot. The robot control system runs directly on the robot's on-board computer, and listens for specific movement commands being sent by the interaction system.

Mimus's people tracking and interaction systems were developed in openFrameworks, a C++-based open source arts-engineering coding toolkit. The robot itself was programmed in RAPID, the proprietary machine language for ABB robots. *Mimus*'s software stack uses following the open-source libraries: openframeworks for graphics, ofxOpencv and ofxCv for image processing, and ofxOsc for inter-application communication over UDP.

PEOPLE TRACKING

Eight-ceiling mounted sensors provide Mimus with a bird's-eye view of her environment. The depth data from these sensors are stitched together into a unified 3D point cloud for the perimeter outside the robot's enclosure. Once the sensors are calibrated to the robot's coordinate system, this depth data enables Mimus to accurately track the head-height of visitors. Although she sees people from above, Mimus can use this tracked head-height to look directly at visitors, face-to-face, as they move around her enclosure (Figure 10).

This people-tracking technique using multiple above-mounted depth sensors was first used in the immersive installation, *Connected Worlds* by Design I/O (2015). This installation used two first-generation Kinects, however the software for *Mimus* expands this technique to integrate up to 12 connected Kinects on one custom-built PC.

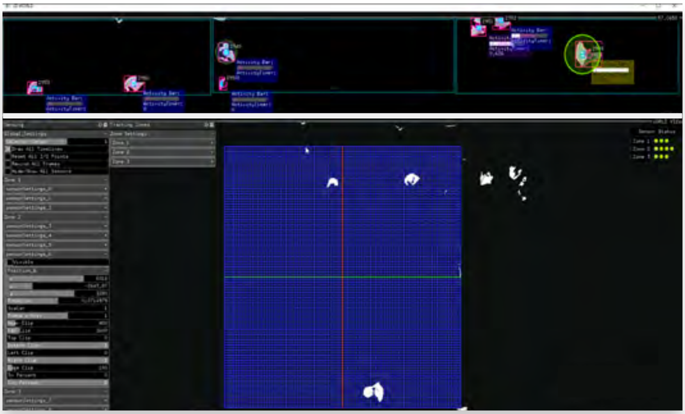

FIGURE 11

People tracking layout: the top horizontal panel provides the robot's CV view; the bottom panel shows the 3D world view.

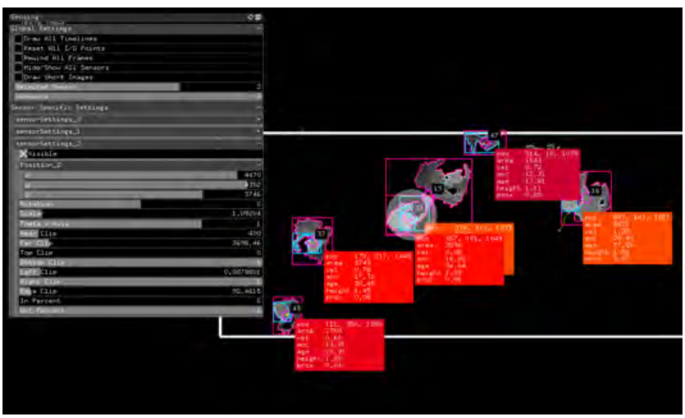

FIGURE 12

Each person detected by the people-tracking system is assigned several quantitive attributes that used to find the most interesting person for Mimus to visit.

The people-tracking software for *Mimus* uses persistent tracking algorithms to detect and follow people-like blobs as they walk around the installation (Figure 11). Each detected person is assigned attributes as they move around the space. This includes their unique ID, real-world position, height, speed, acceleration, area, age, and proximity to robot (Figure 12). These tracked attributes are then used to help the robot decide *who* and *how* to visit a person.

PEOPLE RANKING

Once Mimus can see the people around her, she can then determine who to visit. She uses a ranking algorithm based on peoples tracked attributes to choose the most interesting person in the crowd. This 'most interesting' metric can sometimes be a direct application of tracked attributes: Mimus will prefer people who are *closer*, or *taller*, or *moving the quickest*. Other times, a combination of attributes can be used to infer qualitative relationships from quantitative data. For example, a person who has a high age value and low proximity value has stayed near the robot for a long time: this may infer that some level of trust has developed between the two, and that perhaps Mimus should favor them.

These attributes are dynamically weighted so Mimus can update her metrics for 'most interesting' throughout the day, week, month, or entire exhibit. This adaptive ranking is designed so that people visiting multiple time may have a unique experience with Mimus. For example, one day Mimus may prefer the 'most active' visitors, whereas the next she could prefer to visit the 'shyest'. While on exhibit at the Design Museum London, Mimus predominantly favored people with lower head-heights. This influenced Mimus to prefer children and wheelchair bound visitors before others. Watching children playfully interacting with the robot is a crowd-engaging act, however, this ranking also helped focus Mimus's attention on demographics that are often overlooked in museum exhibition design (Figure 13).

FIGURE 13

Mimus favored visiting shorter people while living at the Design Museum. This had the effect of going to children and wheel-chair bound visitors first.

MOTION PLANNING

Once a person of interest is identified, Mimus begins engaging with them. Mimus's movements are derived from her spatial relationship to other visitors: when a person of interest is close to the enclosure, Mimus come closer to them; if that person moves further away, Mimus retracts back into her space. Furthermore, the scale of movements between visitors and the robot were conveniently correlated with the space of the exhibit. The bodily movements of a person and the robot were approximately the same since the width of the tracking region and the diameter of the robot's reach are equivalent. These mirrored interactions help connect the bodily movements of Mimus to those of visitors. Moreover, the continuous gestural feedback from the robot helped non-expert visitors rapidly internalize the innate kinematic affordances and limitations of specialized machine. However, in early user testing, these 1:1 mirrored movements did not elicit any personable connections from people; the robot's movements were just too robotic. Early users felt like they were operating the robot, instead of interacting with it. Additional interaction design techniques were needed to help the robot transcend its functional form.

Autonomous Behaviors

Three operational strategies were added to Mimus to counteract this predictability and better render life-like qualities on the robot. First, instead of directly mirroring visitor's movements, Mimus integrated an agent-based approach to motion control. Next, Mimus was supplied an attention span so she could become bored with visitors. Finally, the robot could go to sleep when she was tired. These three behaviors gave Mimus enough autonomy to determine her own activities within her enclosure and act more like a living creature. They supplied Mimus with movements that lurched, leaped, and jerked; with motor sounds that chirped, roared, and snored; and with an occasionally indifference for the people around her (Figure 14).

Agent-Based Modeling

The design of Mimus's life-like movements were largely inspired by the autonomous behaviors first explored through *Reverb* in Chapter 3. Like *Reverb*, the virtual geometry underlying Mimus's movements is embedded with its own internal logic for navigating its environment. This agent-based simulation exhibits natural, animal-like movements – like seeking and arriving – when finding a given target point. Other external factors, such as gravity or air resistance, can also added into the virtual environment to elicit different qualities of movement. However, this agent-based motion planning had to overcome many of the innate limitations of industrial robots, including low-latency, real-time control.

The simulated agent and the resulting robot movements often differ in fidelity due to the robot's real-world hardware constraints. To begin, there is noticeable latency between the simulated agent and the robot: the agent runs in a virtual environment at 60 FPS, whereas the robot can only receive data when it is not moving. This drops the robot's data rate to around 4 – 10Hz. Furthermore, the robot must physically maneuver its 1,200kg in space, which takes both time and energy. By contrast, there is no concept of internal mass or inertia factored into the movements of the simulated agent. Despite these differences, the responsive movements of the robot still needed to exhibit life-like characteristics.

In *Mimus*, latency was embraced as an interaction technique instead of an affliction. For example, the time gap between a motion command being and executed was useful for mixing smooth and sharp movements together. This lag would often result in the robot having a run of fluid, serpentine movements that became punctuated with quirky flits and darts. These idiosyncrasies added an element of improvisation and surprise to the robot's otherwise patterned behaviors. Furthermore, the latency from the robot was treated as a delayed reaction time, rather than buffering data: Mimus's slow reflexes facilitated playful, resonate interactions with visitors [Bennett 2015], which became a core user experience with the robot.

Attention Span

Mimus has a tendency towards boredom to solve a few pragmatic problems. First, a single industrial robot can only be oriented to face one person at a time. Limiting her attention span enables Mimus to quickly visit many people and keep the entire crowd engaged. Second, boredom in machines can connote some level of curiosity and intelligence, since becoming bored implies that an opinion has been formulated [Steenson 2017]. This potential for indifference imbues the robot with a life-like, personable characteristic, and requires the active participation of visitors to gain and keep her attention.

Calibrating interactions for Mimus's attention span was a non-trivial task. During early development, I found that if she stayed with a single person too long, others in the crowd would feel ignored and lose interest in the installation. Conversely, if she visited too many people too quickly, visitors would not have enough time to feel a meaningful connection with the robot. Through thorough play testing, I found the acceptable attention span to rang between 4 and 10 seconds per person. Moreover, Mimus's attention span was parameterized to the number of visitors around her: the fewer the visitors, the more time she would spend with each one; the more visitors, the less time. Mimus would spend a minimum of 4 seconds when there were 30+ visitors, and up to 10 seconds when there were only 2 visitors. Visitors

FIGURE 14

Mimus lives in an illuminated glass cage that allows visitors to safely approach the large robot.

FIGURE 15

Mimus goes to sleep when no visitors are near (top) and then wakes up to greet the next person to arrive (bottom).

could also ‘steal’ Mimus’s attention from one another after 4 seconds: if they demonstrate truly obnoxious excitement by frantically waving their arms, then Mimus will leave the current person of interest to seek them out.

Sleep Mode

Mimus goes to sleep once she has visited too many people, and her energy is depleted. The robot enters into *sleep mode* by retracting into a pre-programmed swan-like pose, and then slowly oscillating up and down to visually evoke calm, measured breathing. The sound of breathing is derived from the robot’s motors: the first three motor, which resonate at a deep frequency, are rotated to evoke a sonorous snoring sound. These motors give a low, rhythmic rumble when moved slowly, whereas the upper three motors have a higher pitch when moved.

Sleep mode is the only pre-programmed routine of the interactive installation. However, this routine has two pragmatic purposes. First, it helps set a timescale for visiting the installation: it encourages visitors to move on when the robot is sleeping, and to return when she wakes back up. Secondly, it provides a point of calm contrast for the robot’s lively behaviors and sounds: Mimus’s gentle, rumbling snore (Figure 15a) turns into a sudden roar as her loud motors move to lunge towards a person (Figure 15b). These contrasting movements and sounds were helpful to attract new crowds of people gather and engage with Mimus.

Body Language

Body language is a primitive, yet fluid means of communication. However, it is an underlying means for understanding of the innate behaviors, kinematics and limitations of an unfamiliar thing. Mimus's body language is designed to mimic and compliment that of her visitors. When a person of interest is far away, Mimus is posed higher in the air to orient her machinic gaze down towards them. This posturing was initially implemented for pragmatic reasons: it makes it seem as if the robot is looking over the crowd to a person in the back. However, it had an added effect of making the robot seem quite intimidating and distrusting: like an animal rearing up on its hind legs when confronted. To contrast this sentiment, Mimus takes a warmer, more submissive pose as a person nears: she approaches closely from below, looking upwards toward their face. Her movements also become more jittery when close. This has the effect of making Mimus wiggle with excitement when she face-to-face with a person.

As an installation, Mimus does not try to prescribe a specific emotional response to visitors. Instead, it aims to provide a first point of contact for people to grapple with the complexities of co-existing in a world with intelligent machines. Her poster, movements, and sounds all help the robot transcend its intended mechanical form. Moreover, these physical attributes combine to render into an attentive personality for the robot (Figure 16). Some visitors found Mimus to be flirtatious, playful, or friendly, while others thought the robot to be creepy and were distrusting. In most cases, however, people began to see the robot as more than just a piece of industrial equipment.

Eliciting Life

Mimus deploys several additional strategies related to its physical design, gendering, and end-effector to help visitors see this one-ton lump of metal and motors as a living mechanical creature.

To begin, the appearance of the physical installation is designed to allude to zoos; another familiar space for visiting exotic, and perhaps dangerous creatures. Zoos and menageries provided strong precedents for strategically integrating robust safety infrastructure into engaging audience experiences. They were a valuable design reference for staging a safe space for facilitating awe, wonder, and spectacle with dangerous creatures. Although Mimus is enclosed with standard machine fencing, great care was taken in material selection to foster this desired effect. For example, the enclosure uses 2.2m clear polycarbonate paneling to protect visitors and give uninterrupted visibility to the robot.

Next, Mimus is gendered female to help prime visitors to think of the robot as a creature and not a thing. This decision to gender the robot in all the text about the installation did not come lightly. However, it proved an effective strategy to help visitors better relate and empathize with this piece of industrial equipment as a living thing. Additionally, the name for Mimus alludes to the genus belonging to the mockingbird family, and is derived from the Latin work for ‘mimic’.

The final design decision that helped bring the robot to life was to leave Mimus without any end-effector. Industrial robots are only as useful as tool on the end of them, so leaving a robot unadorned and without a prescribed purpose is unheard of in industry. However, Mimus is simply allowed to exist in the museum, and therefore has no need for any end-effector. When people visit Mimus, they are seeing the robot in her ‘natural’ state: it is how she was birthed from the factory, and how they might find her if she were alive in the wild.

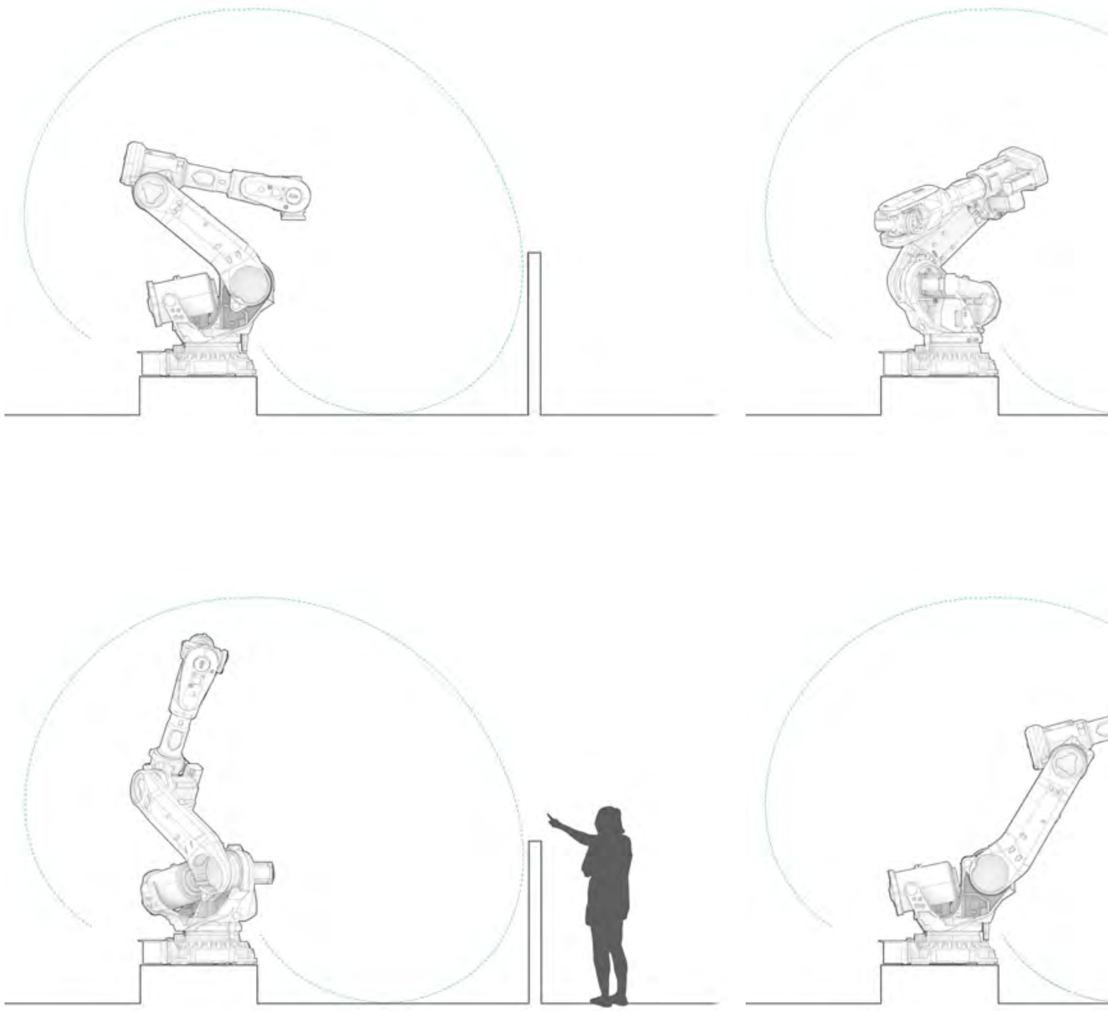

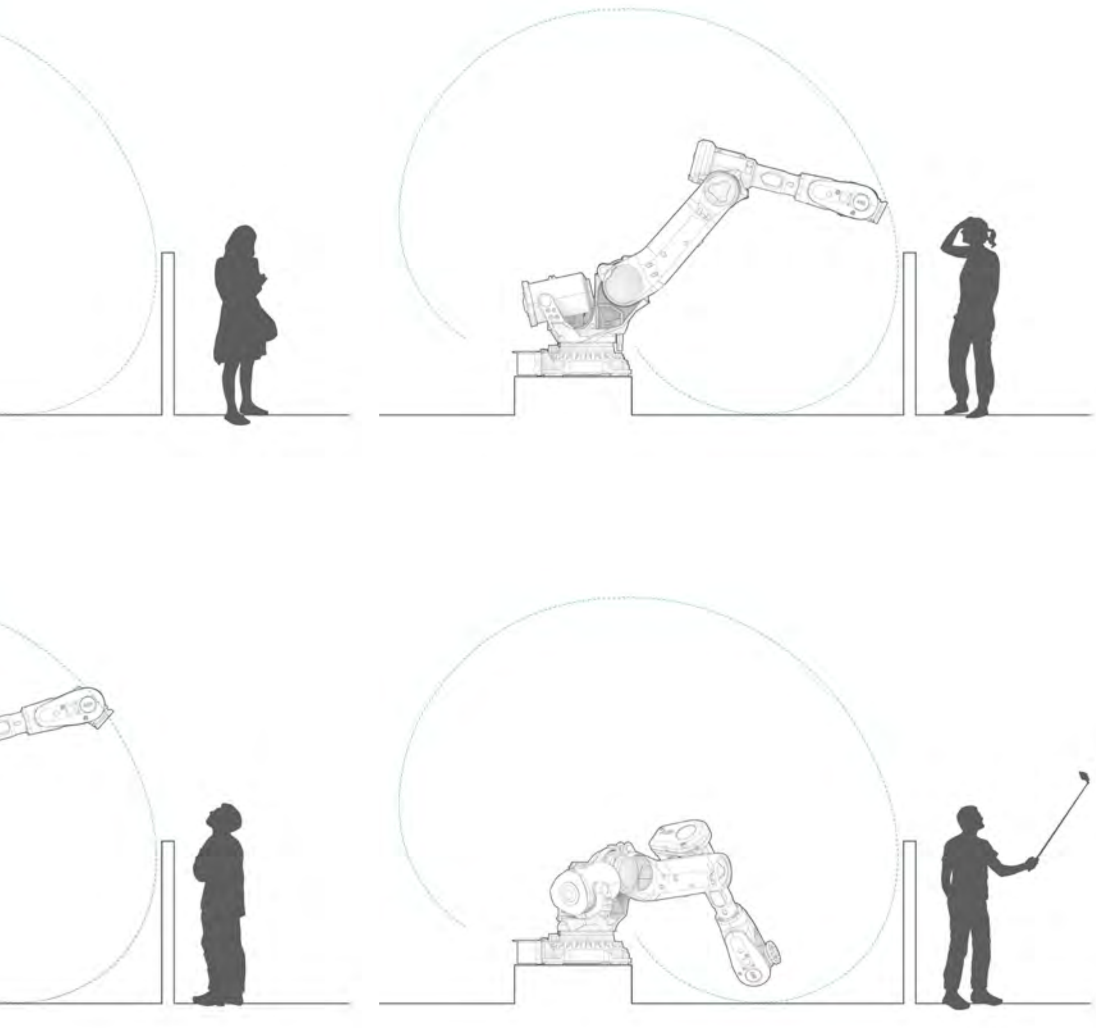

FIGURE 16

These interaction vignettes illustrate different techniques for connecting the body language of visitors and Mimus.

Discussion and Future Work

Mimus offers several pragmatic contributions in human-robot interaction with autonomous machines, including advancing health and safety systems, empathic interactions, and non-verbal communication.

Mimus was a useful case study for bringing industrial robots out of the factory and into live, public environments. To begin, the installation's health and safety systems expanded on existing ISO guidelines for industrial robots. It required much stricter safety standards than specified for traditional manufacturing environments to create a safe and engaging experience for museum visitors and staff. For example, a 200mm floor gap, which is acceptable in factories, was removed to prevent small children from crawling under the paneling. In addition, air gaps between the enclosure's paneling were removed to prevent fingers from getting pinched by the robot. Moreover, the installation operated unsupervised for six months, which prompted additional efforts to child-proof the installation. These unique safety and usability concerns provided the opportunity to identify the necessary adaptations for industrial robots to operate in unstructured environments with crowds of non-specialist people.

Additionally, *Mimus* demonstrates how a non-humanoid robot can effectively engage groups of non-technical users using solely body language and gesture. The interaction design for the installation taps into a person's sense of *pareidolia*: the robot's behaviors broadcast a continuous stream of useful, low-level information on a frequency that humans can't turn off or ignore. Consequently, Mimus feels approachable went she playfully comes close, and intimidating when she looms above your head from afar. These affective bonds can foster gratifying experiences for people, however, they could also serve a tactical purpose. For example, the ability to externally broadcast the internal state of mind of an autonomous robot could be very useful when it is about to something dangerous. This legibility would let the robot produces an affect to instinctually prompt people to step back away from danger.

Mimus pushes the limit of what a non-humanoid robot can communicate with human counterparts: by foregoing traditional human interface devices, this system solely relies on communicative bodily features that are native to the robot. Mimus is limited to three simple outputs to cultivate an empathic bond with visitors: the robot's posture, movements, and sounds. Despite this highly-constrained pallet, Mimus demonstrates the potential for effective human-centered interfaces when thoughtfully designed. Moreover, these interaction techniques could be explored further beyond industrial robotics, and may prove useful for other large, potentially dangerous machines. For example, they may be particularly applicable for autonomous machines – such as aerial drones – which may not be able to integrate traditional UI hardware elements due to cost, computing, or weight limitations.

Finally, Mimus focuses extensively on *body language* as a way of building legibility between a person and an autonomous robot. While this works well for the purposes of an interactive installation, body language and other forms of non-verbal communication do not sufficiently support a wide range of nuanced interactions. Other natural user interfaces, in conjunction with explicit commands and gestures, are a ripe area for further exploration. Natural language interfaces (NLI), paired with hand gestures and body language, is perhaps the ideal combination for fluent communication. Additionally, tangible user interfaces or other computationally augmented tools may prove to be highly intuitive, contextual input devices for working with autonomous machines in domain-specific applications.

Conclusion

This chapter details interaction design strategies for eliciting life-like qualities from a large, potentially dangerous, non-humanoid robot. It begins by sharing my early robotics projects that focus on improving usability and safety for robotics arms in a-typical applications. It then presents *Mimus*, an immersive installation that illustrates how human-centered interfaces can transform an industrial robotic arm into an attentive machinic companion. *Mimus* uses strategic layers of autonomous behaviors to connect the body language of visitors to the body language of the robot. It then discusses the implications of this technique to cultivate empathic bonds between people and machines. While the future may bring additional meaningful ways to communicate with autonomous machines, *Mimus* clearly demonstrates that our existing systems can be reconfigured to foster more empathy and inclusivity with people.

Chapter 7

Conclusion

The interactive systems presented in this dissertation document the progressive embodiment of human-centered interfaces for autonomous fabrication machines. I begin by identifying the fundamental challenges of engaging with machines that can think and make things for themselves. Next, I discuss the most formative existing systems that have influenced my technical and conceptual approach to designing embodied interactions with these machines. I then present three systems – *Reverb*, *Tactum*, and *ExoSkin* – that build on one another to develop interfaces that foster a shared understanding between a designer, an environment, and a fabrication machine. While these systems each bring unique contributions, they also provide the foundational advancements needed to begin exploring meaningful interactions with autonomous fabrication machines.

The final system I present is *Mimus*, a large-scale, immersive installation that transforms a 1,200kg industrial robot into a living, breathing mechanical creature. *Mimus* is the culmination of the technical and conceptual innovations developed throughout this body of research. It demonstrates how an existing, off-the-shelf piece of automation infrastructure can be reconfigured for more human-inclusive purposes. This exemplar human-centered interface introduces the pragmatic and poetic possibilities of truly unmediated interactions with autonomous fabrication machines. In this final chapter, I summarize the key contributions of my dissertation, I reflect on the potential of my research to be extended towards future works, and I discuss its relevancy for investigating future human-machine interactions.

Summary of Contributions

My research has resulted in several contributions relating to Computer-Aided Design/Computer-Aided Manufacturing, Human-Computer Interaction, and Human-Robot Interaction.

To begin, this dissertation has demonstrated several techniques that facilitate new levels of precision for gestural 3D modeling systems. *Reverb* leverages physics simulation to encapsulate the knowledge of an experienced fabricator in its animate, interactive geometry, while also embedding ergonomic design intelligence into its virtual environment. *Tactum* expands on *Reverb*'s fabrication-aware geometry by topologically attaching traditionally modeled 3D geometry to a gesture-manipulated form. Combining these two methods proved critical for 3D modeling around existing objects, and enabled intuitively crafted forms to also have precisely constructed functional parts. These two systems also provide guidelines for intuitive gesture with precise control for gestural 3D modeling systems.

Tactum pushes even further to develop gestural 3D modeling techniques directly on the body. *Tactum* is the *first system* of its kind to focus on *skin* as an interactive canvas for 3D modeling. It presented a unique interface for designing artifacts *in-situ*, at a 1:1 scale in their actual physical context. *ExoSkin* continues this line of inquiry to develop the *first hybrid system* for on-body fabrication. It provides several new workflows for integrating digital techniques into existing high-skill analog craft practices. Although more technical work is required, *ExoSkin* outlines and identifies the unique human, machine, and material challenges that future on-body fabrication systems will need to solve. Moreover, while *Reverb*, *Tactum*, and *ExoSkin* focus specifically on the body, their innovations offer applicable solutions to many other advanced CAD/CAM systems. Lessons from these systems can be applied to any digital fabrication process that needs to operate in dynamically changing environments while working on moving or deformable materials.

Mimus synthesizes the interaction design principles of *Reverb*, *Tactum*, and *ExoSkin* to extrapolate towards unmediated interactions with intelligent, independent fabrication machines. Mimus acts as an important first step into an entirely new field of inquiry: interaction design for autonomous fabrication machines. The ambition of this work is to address basic human-centered design elements for engaging with these autonomous machines. It strives to lay the foundation for creators of future systems to understand the intrinsic benefits challenges of these human-centered interface for autonomous machines.

A brief summary of these contributes is as follows:

Techniques for added precision in gestural 3D modeling.

Techniques for interactive fabrication-aware geometry.

A skin-based approach to 3D modeling.

A framework for designing on-body CAD systems.

Techniques for on-body fabrication of wearable objects.

A framework for designing on-body fabrications systems.

Two open-source libraries for interacting with robotic arms.

Interaction design techniques for autonomous robotic arms.

Outcomes

I am humbled by the enthusiastic external interest and support for my work throughout my tenure as a PhD candidate. I am privileged to have had my research internationally exhibited at leading cultural institutions, published at ACM conferences, and widely covered by diverse media outlets across design, art, and technology communities.

Over the last three years, over 35 academic papers related to wearables, on-body design, or on-body fabrication have referenced my body of work. I am especially proud of Tactum, which also received a Best Paper Nominee award at ACM CHI conference in 2015. However, this body of work has also made a notable impact outside of academia. Reverb was exhibited as a part of a 3D printed fashion event for New York Fashion Week in 2013. Tactum was featured in *The Body Engaged* exhibit by the Boston Design Museum in 2015.

Perhaps my most ambitious and challenging achievement has been *Mimus*: a commission from the Design Museum London for their inaugural exhibit, *Fear and Love*. I am proud of the impact that *Mimus* has had in conveying the importance of interaction design for robotics to the global Design community. My work in this has been recognized as one of the 10 Biggest Design Trends of 2017 by *Dezeen* (the world's leading design news website, at this time of writing). *Mimus* also received an Ars Electronica *STARTS Prize* Honorable Mention, and was central to me being named a World Economic Forum Cultural Leader in 2017.

My research has also been widely covered in a variety of media outlets, such as the BBC, the Guardian, Discover, Wired, Popular Science, Vice Media, and numerous others. While this external validation is certainly confidence-boosting, it is a valuable form of feedback to continuously improve the clarity, quality, and communication of my work. While the technical implementations of my research will be frozen in a moment in time, I hope the concepts and design strategies I have demonstrated can continue to contribute to the next wave of academic and cultural inquiries into the future of human-machine interaction.

Future Work

The main thrust of this research validates the potential for non-experts to have meaningful engagements with highly technical fabrication machines: *Mimus* demonstrates that it is possible; *Reverb*, *Tactum*, and *ExoSkin* detail the techniques and frameworks that make it so. Some interfaces I present here show examples of large, dangerous machines operating inches away from a person. While these more intelligent, contextual interfaces demonstrate that advanced manufacturing can be tamed for more human-inclusive purposes, these demonstrations were developed in highly controlled environments for highly constrained use cases. Additional technical advancements in robotic perception and actuation are needed before these research prototypes could become reliably and safely adopted for real-world scenarios.

Techniques in artificial intelligence and machine learning also show promise for improving how these autonomous systems can dynamically respond to changing contexts. In this research, I rely on cybernetic models of intelligence to elicit life-like behaviors from these interactive systems. This strategy is effective for driving interactions for geometries or machines with limited means of external expression. However, it does not adequately capture more subtle forms of input from people. Each system presented in this dissertation is limited in the nuance and dexterity of the natural gestures it detects and tracks. Future work focusing on more sophisticated artificial intelligence and machine learning techniques may provide fabrication machines a more innate and holistic understanding of their human counterparts.

More intelligent, reliable, and safer human-centered interfaces open the door to exploring new and novel domain spaces for autonomous fabrication machines. *Mimus*, for example, showcases how an autonomous fabrication machine might live in a public setting. However, the functionality of the robot is limited: Mimus does not have a purpose or use beyond her own existence. Many of its unique abilities – its strength, precision, and adaptability, for example – are not fully explored in this context. Future work could focus on leveraging the unique affordances of other fabrication machines outside of traditional manufacturing applications. Adapting these machines to live in laboratories, on film sets, in public spaces, or in domestic settings has the potential to spur diverse and compelling alternatives for how we can more meaningfully co-exist.

This research could also be expanded in new, fruitful directions by exploring the unique affordances of human-centered interaction with a team of autonomous machines. My body of work examines ways of creating empathic connections between a machine and a single-user, multiple users, or a crowd of people. However, it would be interesting to investigate how the dynamic between a person and machine changes once the machines outnumber the person. Furthermore, developing interfaces that combine the unique abilities of a person within an entire ecosystem of robots is an equally exciting territory. For example, enabling a drone, robotic arm, quadruped robot, snake robot, and a human to all fluidly communicate with one another could unlock an amazing set of yet-to-be-explored abilities.

Finally, many of the design decisions for human-centered interfaces with autonomous systems come down to negotiating the right balance of human agency and machine autonomy. In my systems, these design decisions are made through intuition, experience, and trail-and-error. Translating this trail-and-error process into a standard set of generalizable principles is a challenging problem. It requires an extensive, high-dimensional survey of all the permutations and variations of different autonomous machines doing different tasks in different environments to truly map out the design space. While unfortunately out of the scope of this dissertation, formalizing this intuitive knowledge into a cohesive design language for crafting interactions with autonomous machines could provide an impactful advancement towards solidifying a framework for human-robot interaction.

Final Thoughts

There is a fine line between automation that helpful or harmful to people, and it can be difficult to gauge whether an autonomous system is replacing or expanding a person's capabilities. The metric I have used in my own systems asks whether a person and a machine are creating something together that neither could do apart: if the machine makes things slightly more convenient, or if a group of people could do it better, then this human-machine system is not ambitious enough. This method of reflective audit has proved very useful for thinking beyond the most common use cases for human-robot interaction. Integrating this line of inquiry during the design process is helpful for escaping our most common preconceptions for these machines.

Today, the prevailing purpose for a robot – intelligent or otherwise – is to function as a servant: the more complicated task it can achieve, the better the robot. However, this lacks imagination for any potential to enhance our lives in more meaningful ways. *Mimus* aimed to demonstrate how an industrial robot – a tool of automation – could be transformed into an attentive companion. Eliciting life from this machine brought a new purpose to an old tool; enabling it to adore, amaze, astound, and entertain a crowd. As these machines become increasingly intelligent, we will need more compelling examples of alternative, symbiotic human-machine relationships. While *Mimus* was a useful first step, my hope is that this work can ignite a wider range of interest for exploring a-typical affordances with these machines.

My research is driven by a relentless quest to rediscover valuable ideas that have been overlooked or underexplored in the digital designing and making of things. The diverse forms of output that my systems produce reveal my roaming curiosity: *Reverb* explores 3D printed fashion, *Tactum* and *ExoSkin* examine wearables, and much of my work, both formal and informal, has revolved around robotics. However, the importance of these systems is not what they produce, but how they engage and empower. I have shown how digital design tools can extend our own creativity, how digital interfaces can engage their physical contexts, how fabrication machines can work in close synchronicity with a human counterpart, and how life-like, personable qualities can be successfully rendered onto giant, non-humanoid machines. What connects all these areas of inquiry is the desire to appropriate automation in ways that expand, augment, and enhance human capabilities.

This dissertation documents a journey that begins with solving pragmatic problems in CAD/CAM and ends at the threshold of a new territory in human-machine interaction. As autonomous machines become a more ubiquitous part of everyday life, it is critical that we have more effective ways of interacting and communicating with one another. Moreover, deciding how we will coexist with intelligent, attentive, and animate machines is a timely and relevant area of inquiry, given our contentious relationship with automation today. While the negative side-effects of automation may seem like daunting forces to overcome, we are not bystanders to technological change. If this body of work offers nothing more, I hope it conveys a sense of optimism and empowerment for redirecting these autonomous systems towards more inclusive, human-centered alternatives.

Bibliography

1. *Fraser Anderson, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2013. YouMove: enhancing movement training with an augmented reality mirror. In Proceedings of the 26th annual* ACM *symposium on User interface software and technology* (UIST '13). ACM, *New York*, NY, USA, 311-320. DOI: *http://dx.doi.org/10.1145/2501988.2502045*

2. *Daniel Ashbrook, Shitao Stan Guo, and Alan Lambie. 2016. Towards Augmented Fabrication: Combining Fabricated and Existing Objects. In Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (CHI EA '16). ACM, *New York*, NY, USA, 1510-1518. DOI: *https://doi.org/10.1145/2851581.2892509*

3. *Joshua Bard, Madeline Gannon, Zachary Jacobson-Weaver, Michael Jeffers, Brian Smith, and Mauricio Contreras (2014) Seeing is Doing: Synthetic Tools for Robotically Augmented Fabrication in High-Skill Domains.* ACADIA 14: *Design Agency. In Proceedings of the 34th Annual Conference of the Association for Computer Aided Design in Architecture* (ACADIA) Los Angeles, USA: *pp.* 409-416.

4. *Peter Bennett, Stuart Nolan, Ved Uttamchandani, Michael Pages, Kirsten Cater, and Mike Fraser. 2015. Resonant Bits: Harmonic Interaction with Virtual Pendulums. In Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction* (TEI '15). ACM, *New York*, NY, USA, 49-52. DOI: *https://doi.org/10.1145/2677199.2680569*

5. *Pierre Bézier. 1986. The Mathematical Basis of the* UNISURF CAD *System. Butterworth-Heinemann, Newton*, MA, USA.

6. *Richard A. Bolt. 1980. "Put-that-there": Voice and gesture at the graphics interface. In Proceedings of the 7th annual conference on Computer graphics and interactive techniques* (SIGGRAPH '80). ACM, *New York*, NY, USA, 262-270. DOI=*http://dx.doi.org/10.1145/800250.807503*

7. *Cynthia Breazeal Brian Scassellati, "How to build robots that make friends and influence people," In Proceedings* 1999 IEEE/RSJ *International Conference on Intelligent Robots and Systems* (IROS '99). 858-863 *vol.2.* DOI=10.1109/IROS.1999.812787

8. *Erik Brynjolfsson and Andrew McAfee. The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies. First Edition. New York:* W. W. *Norton & Company*, 2014.

9. *Liwei Chan, Rong-Hao Liang, Ming-Chang Tsai, Kai-Yin Cheng, Chao-Huai Su, Mike* Y. *Chen, Wen-Huang Cheng, and Bing-Yu Chen. 2013. FingerPad: private and subtle interaction using fingertips. In Proceedings of the 26th annual* ACM *symposium on User interface software and technology* (UIST '13). ACM, *New York*, NY, USA, 255-260. DOI: *http://dx.doi.org/10.1145/2501988.2502016*

10. *Ke-Yu Chen, Kent Lyons, Sean White, and Shwetak Patel. 2013. uTrack: 3D input using two magnetic sensors. In Proceedings of the 26th annual* ACM *symposium on User interface software and technology* (UIST '13). ACM, *New York*, NY, USA, 237-244. DOI: *http://dx.doi.org/10.1145/2501988.2502035*

11. *Steven A. Coons. An Outline of the Requirements for a Computer-Aided Design System.* M.I.T. *Electronic Systems Laboratory. Technical Memorandum* ESL-TM-169. *Cambridge:* M.I.T. *Electronic Systems Laboratory*, 1963.

12. *Stelian Coros, Bernhard Thomaszewski, Gioacchino Noris, Shinjiro Sueda, Moira Forberg, Robert W. Sumner, Wojciech Matusik, and Bernd Bickel. 2013. Computational design of mechanical characters.* ACM *Trans. Graph.* 32, 4, *Article* 83 *(July 2013), 12 pages.* DOI: *https://doi.org/10.1145/2461912.2461953*

13. *Kerstin Dautenhahn. 1997. Ants Don't Have Friends — Thoughts on Socially Intelligent Agents. In Socially Intelligent Agents, pages* 22–27. AAAI *Press, Technical report* FS-97-02, 1997.

14. *Kerstin Dautenhahn. 1999. Embodiment and interaction in socially intelligent life-like agents. In Computation for metaphors, analogy, and agents, Chrystopher L. Nehaniv (Ed.). Lecture Notes in Computer Science, Vol. 1562. Springer-Verlag, Berlin, Heidelberg* 102-141.

15. *Laura Devendorf and Kimiko Ryokai. 2015. Being the Machine: Reconfiguring Agency and Control in Hybrid Fabrication. In Proceedings of the 33rd Annual* ACM *Conference on Human Factors in Computing Systems* (CHI '15). ACM, *New York*, NY, USA, 2477-2486. *http://doi.acm.org/10.1145/2702123.2702547*

16. *Donald M. Fisk. American Labor in the 20th Century. Compensation and Working Conditions. Fall 2001. Bureau of Labor Statistics.*

17. *Sean Follmer, David Carr, Emily Lovell, and Hiroshi Ishii. 2010. CopyCAD: remixing physical objects with copy and paste from the real world. In Adjunct proceedings of the 23nd annual* ACM *symposium on User interface software and technology* (UIST '10). ACM, *New York*, NY, USA, 381-382. DOI: *https://doi.org/10.1145/1866218.1866230*

18. *Sean Follmer and Hiroshi Ishii. 2012. KidCAD: digitally remixing toys through tangible tools. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12). ACM, New York, NY, USA, 2401-2410. DOI: http://dx.doi.org/10.1145/2207676.2208403*

19. *Madeline Gannon. 2016. Mimus: Coming Face-to-Face With Our Companion Species. In "Fear And Love: Reactions to a Complex World". McGuirk, J., and Herrero, G, (eds.) Phaidon Press, Ltd. London, UK.*

20. *Madeline Gannon. 2016. Open-Source Robotics. In Openism: Conversations on Open Hardware. Newman, A., Tarasiewicz, M., Wagner, S.C., Wuschitz, S. (eds.). University of Applied Arts Vienna. Vienna, Austria.*

21. *Madeline Gannon. The Shape of Touch: On-Body Interfaces for Digital Design and Fabrication. Architectural Design, 87:6. (2017). 114–119. DOI=10.1002/ad.2246*

22. *Madeline Gannon, Tovi Grossman, and George Fitzmaurice. 2015. Tactum: A Skin-Centric Approach to Digital Design and Fabrication. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15). ACM, New York, NY, USA, 1779-1788. DOI: http://doi.acm.org/10.1145/2702123.2702581*

23. *Madeline Gannon, Tovi Grossman, and George Fitzmaurice. 2016. ExoSkin: On-Body Fabrication. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16). ACM, New York, NY, USA, 5996-6007. DOI: https://doi.org/10.1145/2858036.2858576*

24. *Neil Gershenfeld. 2007. Fab: The Coming Revolution on Your Desktop--From Personal Computers to Personal Fabrication. Basic Books, Inc., New York, NY, USA.*

25. *Chris Harrison, Desney Tan, and Dan Morris. 2010. Skinput: appropriating the body as an input surface. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10). ACM, New York, NY, USA, 453-462. DOI=10.1145/1753326.1753394.*

26. *Chris Harrison, Hrvoje Benko, and Andrew D. Wilson. 2011. OmniTouch: wearable multitouch interaction everywhere. In Proceedings of the 24th annual ACM symposium on User interface software and technology (UIST '11). ACM, New York, NY, USA, 441-450. http://doi.acm.org/10.1145/2047196.2047255*

27. *Mark D Gross and Ariel Kemp. Gesture Modeling: Using Video to Capture Freehand Modeling Commands. CAAD Futures '01, 33-46.*

28. *Chris Harrison, Desney Tan, and Dan Morris. 2011. Skinput: appropriating the skin as an interactive canvas. Commun. ACM 54, 8 (August 2011), 111-118. DOI: https://doi.org/10.1145/1978542.1978564*

29. *Chris Harrison, Hrvoje Benko, and Andrew D. Wilson. 2011. OmniTouch: wearable multitouch interaction everywhere. In Proceedings of the 24th annual ACM symposium on User interface software and technology (UIST '11). ACM, New York, NY, USA, 441-450. DOI: https://doi.org/10.1145/2047196.2047255*

30. *Chris Harrison, Shilpa Ramamurthy, and Scott E. Hudson. 2012. On-body interaction: armed and dangerous. In Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction (TEI '12), Stephen N. Spencer (Ed.). ACM, New York, NY, USA, 69-76. DOI: https://doi.org/10.1145/2148131.2148148*

31. *Ke Huo, Vinayak, and Karthik Ramani. 2017. Window-Shaping: 3D Design Ideation by Creating on, Borrowing from, and Looking at the Physical World. In Proceedings of the Eleventh International Conference on Tangible, Embedded, and Embodied Interaction (TEI '17). ACM, New York, NY, USA, 37-45. DOI: https://doi.org/10.1145/3024969.3024995*

32. *Takeo Igarashi, Satoshi Matsuoka, Sachiko Kawachiya, and Hidehiko Tanaka. 1997. Interactive beautification: a technique for rapid geometric design. In Proceedings of the 10th annual ACM symposium on User interface software and technology (UIST '97). ACM, New York, NY, USA, 105-114. DOI=10.1145/263407.263525 http://doi.acm.org/10.1145/263407.263525*

33. *Juan David Hincapié-Ramos, Xiang Guo, Paymahn Moghadasian, and Pourang Irani. 2014. Consumed endurance: a metric to quantify arm fatigue of mid-air interactions. In Proceedings of the 32nd annual ACM conference on Human factors in computing systems (CHI '14). ACM, New York, NY, USA, 1063-1072. DOI=10.1145/2556288.2557130 http://doi.acm.org/10.1145/2556288.2557130*

34. *Ken Hinckley, Randy Pausch, John C. Goble, and Neal F. Kassell. 1994. A survey of design issues in spatial input. In Proceedings of the 7th annual ACM symposium on User interface software and technology (UIST '94). ACM, New York, NY, USA, 213-222. DOI=http://dx.doi.org/10.1145/192426.192501*

35. Christian Holz and Andrew Wilson. 2011. Data miming: inferring spatial object descriptions from human gesture. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11). ACM, New York, NY, USA, 811–820. DOI: https://doi.org/10.1145/1978942.1979060

36. Christian Holz, Tovi Grossman, George Fitzmaurice, and Anne Agur. 2012. Implanted user interfaces. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12). ACM, New York, NY, USA, 503-512. DOI: http://dx.doi.org/10.1145/2207676.2207745

37. Gabe Johnson, Mark Gross, Ellen Yi-Luen Do, and Jason Hong. 2012. Sketch it, make it: sketching precise drawings for laser cutting. In CHI '12 Extended Abstracts on Human Factors in Computing Systems (CHI EA '12). ACM, New York, NY, USA, 1079-1082. DOI: http://dx.doi.org/10.1145/2212776.2212390

38. Hyosun Kim, Georgia Albuquerque, Sven Havemann, and Dieter W. Fellner. 2005. Tangible 3D: hand gesture interaction for immersive 3D modeling. In Proceedings of the 11th Eurographics conference on Virtual Environments (EGVE'05), Erik Kjems and Roland Blach (Eds.). Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 191–199. DOI=http://dx.doi.org/10.2312/EGVE/IPT_EGVE2005/191-199

39. Yongkwan Kim and Seok-Hyung Bae. 2016. SketchingWithHands: 3D Sketching Handheld Products with First-Person Hand Posture. In Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16). ACM, New York, NY, USA, 797-808. DOI: https://doi.org/10.1145/2984511.2984567

40. Heather Knight. How Humans Respond to Robots: Building Public Policy through Good Design. The Robots Are Coming: The Project On Civilian Robotics. July 2014. Brookings Institute.

41. Myron W. Krueger. 1977. Responsive environments. In Proceedings of the June 13-16, 1977, national computer conference (AFIPS '77). ACM, New York, NY, USA, 423-433. DOI=http://dx.doi.org/10.1145/1499402.1499476

42. Manfred Lau, Masaki Hirose, Akira Ohgawara, Jun Mitani, and Takeo Igarashi. 2012. Situated modeling: a shape-stamping interface with tangible primitives. In Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction (TEI '12), Stephen N. Spencer (Ed.). ACM, New York, NY, USA, 275-282. DOI: https://doi.org/10.1145/2148131.2148190

43. Daniel Leithinger, David Lakatos, Anthony DeVincenzi, Matthew Blackshaw, and Hiroshi Ishii. 2011. Direct and gestural interaction with relief: a 2.5D shape display. In Proceedings of the 24th annual ACM symposium on User interface software and technology (UIST '11). ACM, New York, NY, USA, 541–548. DOI: https://doi.org/10.1145/2047196.2047268

44. Shu-Yang Lin, Chao-Huai Su, Kai-Yin Cheng, Rong-Hao Liang, Tzu-Hao Kuo, and Bing-Yu Chen. 2011. Pub - point upon body: exploring eyes-free interaction and methods on an arm. In Proceedings of the 24th annual ACM symposium on User interface software and technology (UIST '11). ACM, New York, NY, USA, 481–488. DOI: https://doi.org/10.1145/2047196.2047259

45. Hsin-Liu (Cindy) Kao, Christian Holz, Asta Roseway, Andres Calvo, and Chris Schmandt. 2016. DuoSkin: rapidly prototyping on-skin user interfaces using skin-friendly materials. In Proceedings of the 2016 ACM International Symposium on Wearable Computers (ISWC '16). ACM, New York, NY, USA, 16-23. DOI: https://doi.org/10.1145/2971763.2971777

46. Joanne Lo, Doris Jung Lin Lee, Nathan Wong, David Bui, and Eric Paulos. 2016. Skintillates: Designing and Creating Epidermal Interactions. In Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16). ACM, New York, NY, USA, 853-864. DOI: https://doi.org/10.1145/2901790.2901885

47. Ignacio Llamas, Byungmoon Kim, Joshua Gargus, Jarek Rossignac, and Chris D. Shaw. 2003. Twister: a space-warp operator for the two-handed editing of 3D shapes. ACM Trans. Graph. 22, 3 (July 2003), 663-668. DOI: https://doi.org/10.1145/882262.882323

48. Marshall McLuhan. (1964). Understanding media: The extensions of man ([2d ed.].). New York: New American Library.

49. Roger K. Moore. A Bayesian explanation of the "Uncanny Valley" effect and related psychological phenomena. Scientific Reports, 2, 864 (2012). http://doi.org/10.1038/srep00864

50. Masahiro Mori. 1970. Bukimi no tani {The uncanny valley}. Energy. v7 i4. 33-35.

51. Yuki Mori and Takeo Igarashi. 2007. Plushie: an interactive design system for plush toys. ACM Trans. Graph. 26, 3, Article 45 (July 2007). DOI: https://doi.org/10.1145/1276377.1276433

52. *Stefanie Mueller, Pedro Lopes, and Patrick Baudisch. 2012. Interactive construction: interactive fabrication of functional mechanical devices. In Proceedings of the 25th annual ACM symposium on User interface software and technology (UIST '12). ACM, New York, NY, USA, 599-606. DOI = 10.1145/2380116.2380191*

53. *Tao Ni, Amy K. Karlson, and Daniel Wigdor. 2011. AnatOnMe: facilitating doctor-patient communication using a projection-based handheld device. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11). ACM, New York, NY, USA, 3333-3342. http://doi.acm.org/10.1145/1978942.1979437*

54. *Aditya Shekhar Nittala and Jürgen Steimle. 2016. Digital fabrication pipeline for on-body sensors: design goals and challenges. In Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct (UbiComp '16). ACM, New York, NY, USA, 950-953. DOI: https://doi.org/10.1145/2968219.2979140*

55. *Noble, David F. Forces of Production: A Social History of Industrial Automation. New York: Knopf, 1984.*

56. *Masa Ogata, Yuta Sugiura, Yasutoshi Makino, Masahiko Inami, and Michita Imai. 2013. SenSkin: adapting skin as a soft interface. In Proceedings of the 26th annual ACM symposium on User interface software and technology (UIST '13). ACM, New York, NY, USA, 539-544. DOI: http://dx.doi.org/10.1145/2501988.2502039*

57. *Simon Olberding, Kian Peen Yeo, Suranga Nanayakkara, and Jurgen Steimle. 2013. AugmentedForearm: exploring the design space of a display-enhanced forearm. In Proceedings of the 4th Augmented Human International Conference (AH '13). ACM, New York, NY, USA, 9-12. DOI=http://dx.doi.org/10.1145/2459236.2459239*

58. *Huaishu Peng, Amit Zoran, and François V. Guimbretière. 2015. D-Coil: A Hands-on Approach to Digital 3D Models Design. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15). ACM, New York, NY, USA, 1807-1815. DOI=10.1145/2702123.2702381*

59. *Pew Research Center. AI, Robotics, and the Future of Jobs. Digital Life in 2025 Report. August 2014.*

60. *Piroozfar, Poorang A. E., and Frank T. Piller. Mass Customisation and Personalisation in Architecture and Construction. New York: Routledge, 2013.*

61. *Helmut Pottmann. Architectural Geometry and Fabrication-Aware Design. Nexus Network Journal, 15:2, August 2013, Springer, Basel. 195-208.*

62. *Alec Rivers, Ilan E. Moyer, and Frédo Durand. 2012. Position-correcting tools for 2D digital fabrication. ACM Trans. Graph. 31, 4, Article 88 (July 2012), 7 pages. http://doi.acm.org/10.1145/2185520.2185584*

63. *Rajinder Sodhi, Hrvoje Benko, and Andrew Wilson. 2012. LightGuide: projected visualizations for hand movement guidance. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12). ACM, New York, NY, USA, 179-188. http://doi.acm.org/10.1145/2207676.2207702*

64. *Jaisa Reichardt, Traub, J. F. 1., McCorduck, P., & Institute of Contemporary Arts (London). 1969. Cybernetic serendipity: The computer and the arts. New York: Praeger.*

65. *Brian Reffin Smith. Soft Computing: Art and Design, Addison-Wesley, 1984, 147–155.*

66. *Craig W. Reynolds. 1987. Flocks, herds and schools: A distributed behavioral model. In Proceedings of the 14th annual conference on Computer graphics and interactive techniques (SIGGRAPH '87), Maureen C. Stone (Ed.). ACM, New York, NY, USA, 25-34. DOI: https://doi.org/10.1145/37401.37406*

67. *Douglas T. Ross and Jorge E. Rodriguez. 1963. Theoretical foundations for the computer-aided design system. In Proceedings of the May 21-23, 1963, spring joint computer conference (AFIPS '63 (Spring)). ACM, New York, NY, USA, 305-322. DOI=http://dx.doi.org/10.1145/1461551.1461589*

68. *Daniel Saakes, Hui-Shyong Yeo, Seung-Tak Noh, Gyeol Han, and Woontack Woo. 2016. Mirror Mirror: An On-Body T-shirt Design System. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI ‹16). ACM, New York, NY, USA, 6058-6063. DOI: https://doi.org/10.1145/2858036.2858282*

69. *Greg Saul, Manfred Lau, Jun Mitani, and Takeo Igarashi. 2010. SketchChair: an all-in-one chair design system for end users. In Proceedings of the fifth international conference on Tangible, embedded, and embodied interaction (TEI ‹11). ACM, New York, NY, USA, 73-80. DOI=http://dx.doi.org/10.1145/1935701.1935717*

70. *Steven Schkolne, Michael Pruett, and Peter Schröder. 2001. Surface drawing: creating organic 3D shapes with the hand and tangible tools. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01). ACM, New York, NY, USA, 261-268. DOI=http://dx.doi.org/10.1145/365024.365114*

71. *Jia Sheng, Ravin Balakrishnan, and Karan Singh. 2006. An interface for virtual 3D sculpting via physical proxy. In Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia (GRAPHITE '06). ACM, New York, NY, USA, 213-220. DOI=http://dx.doi.org/10.1145/1174429.1174467*

72. *Roy Shilkrot, Pattie Maes, Joseph A. Paradiso, and Amit Zoran. 2015. Augmented Airbrush for Computer Aided Painting (CAP). ACM Trans. Graph. 34, 2, Article 19 (March 2015), 11 pages. DOI: https://doi.org/10.1145/2699649*

73. *Ross T. Smith, Bruce H. Thomas, and Wayne Piekarski. 2008. Digital foam interaction techniques for 3D modeling. In Proceedings of the 2008 ACM symposium on Virtual reality software and technology (VRST '08). ACM, New York, NY, USA, 61-68. DOI: https://doi.org/10.1145/1450579.1450592*

74. *Molly Wright Steenson. Architectural Intelligence: How Designers and Architects Created the Digital Landscape. MIT Press. 2017.*

75. *Ivan E. Sutherland. 1963. Sketchpad: a man-machine graphical communication system. In Proceedings of the May 21-23, 1963, spring joint computer conference (AFIPS '63 (Spring)). ACM, New York, NY, USA, 329-346. DOI=http://dx.doi.org/10.1145/1461551.1461591*

76. *Leila Takayama and Caroline Pantofaru. 2009. Influences on proxemic behaviors in human-robot interaction. In Proceedings of the 2009 IEEE/RSJ international conference on Intelligent robots and systems (IROS'09). IEEE Press, Piscataway, NJ, USA, 5495-5502.*

77. *Alexander Teibrich, Stefanie Mueller, François Guimbretière, Robert Kovacs, Stefan Neubert, and Patrick Baudisch. 2015. Patching Physical Objects. In Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology (UIST '15). ACM, New York, NY, USA, 83-91. DOI=http://dx.doi.org/10.1145/2807442.2807467*

78. *Nobuyuki Umetani, Danny M. Kaufman, Takeo Igarashi, and Eitan Grinspun. 2011. Sensitive couture for interactive garment modeling and editing. In ACM SIGGRAPH 2011 papers (SIGGRAPH '11), Hugues Hoppe (Ed.). ACM, New York, NY, USA, Article 90, 12 pages. DOI: https://doi.org/10.1145/1964921.1964985*

79. *Pascal Volino, Frederic Cordier, and Nadia Magnenat-Thalmann. 2005. From early virtual garment simulation to interactive fashion design. Comput. Aided Des. 37, 6 (May 2005), 593-608. DOI=http://dx.doi.org/10.1016/j.cad.2004.09.003*

80. *Jin Wang, Guodong Lu, Weilong Li, Long Chen, and Yoshiyuki Sakaguti. 2009. Interactive 3D garment design with constrained contour curves and style curves. Comput. Aided Des. 41, 9 (September 2009), 614-625. DOI=http://dx.doi.org/10.1016/j.cad.2009.04.009*

81. *Robert Wang, Sylvain Paris, and Jovan Popović. 2011. 6D hands: markerless hand-tracking for computer aided design. In Proceedings of the 24th annual ACM symposium on User interface software and technology (UIST '11). ACM, New York, NY, USA, 549-558. DOI: https://doi.org/10.1145/2047196.2047269*

82. *Martin Weigel, Vikram Mehta, and Jürgen Steimle. 2014. More than touch: understanding how people use skin as an input surface for mobile computing. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14). ACM, New York, NY, USA, 179-188. DOI: https://doi.org/10.1145/2556288.2557239*

83. *Martin Weigel, Tong Lu, Gilles Bailly, Antti Oulasvirta, Carmel Majidi, and Jürgen Steimle. 2015. iSkin: Flexible, Stretchable and Visually Customizable On-Body Touch Sensors for Mobile Computing. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15). ACM, New York, NY, USA, 2991-3000. DOI: https://doi.org/10.1145/2702123.2702391*

84. *Christian Weichel, Manfred Lau, David Kim, Nicolas Villar, and Hans W. Gellersen. 2014. MixFab: a mixed-reality environment for personal fabrication. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14). ACM, New York, NY, USA, 3855-3864. DOI: https://doi.org/10.1145/2556288.2557090*

85. *Amy Wibowo, Daisuke Sakamoto, Jun Mitani, and Takeo Igarashi. 2012. DressUp: a 3D interface for clothing design with a physical mannequin. In Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction (TEI '12), Stephen N. Spencer (Ed.). ACM, New York, NY, USA, 99-102. http://doi.acm.org/10.1145/2148131.2148153*

86. *Norbert Wiener. Cybernetics: Or, Control and Communication in the Animal and the Machine. 2d ed. New York: M.I.T. Press, 1961.*

87. *World Economic Forum. The Future of Jobs: Employment, Skills and Workforce Strategy for the Fourth Industrial Revolution. Global Challenge Insights Report. January 2016.*

88. *World Economic Forum. Advancing Human-Centred Economic Progress in the Fourth Industrial Revolution. G20/T20 Policy Brief. May 2017.*

89. *Ian D. Wyatt and Daniel E. Hecker. Occupational Changes During the 20th Century. Monthly Labor Review. March 2005. U.S. Bureau of Labor Statistics.*

90. *Yamashita, M. M., Yamaoka, J., and Kakehi, Y. Enchanted Scissors: a scissor interface for support in cutting and interactive fabrication. In ACM SIGGRAPH 2013 Posters. SIGGRAPH '13.*

91. *Xing-Dong Yang, Tovi Grossman, Daniel Wigdor, and George Fitzmaurice. 2012. Magic finger: always-available input through finger instrumentation. In Proceedings of the 25th annual ACM symposium on User interface software and technology (UIST '12). ACM, New York, NY, USA, 147-156. DOI=10.1145/2380116.2380137.*

92. *Yupeng Zhang, Teng Han, Zhimin Ren, Nobuyuki Umetani, Xin Tong, Yang Liu, Takaaki Shiratori, and Xiang Cao. 2013. BodyAvatar: creating freeform 3D avatars using first-person body gestures. In Proceedings of the 26th annual ACM symposium on User interface software and technology (UIST '13). ACM, New York, NY, USA, 387-396. DOI: http://dx.doi.org/10.1145/2501988.2502015*

93. *Aleksandar Zivanovic. 2005. The development of a cybernetic sculptor: Edward Ihnatowicz and the Schrenster. In Proceedings of the 5th conference on Creativity & cognition (C&C ‹05). ACM, New York, NY, USA, 102-108. DOI=http://dx.doi.org/10.1145/1056224.1056240*

94. *Kening Zhu, Alexandru Dancu, and Shengdong (Shen) Zhao. 2016. FusePrint: A DIY 2.5D Printing Technique Embracing Everyday Artifacts. In Proceedings of the 2016 ACM Conference on Designing Interactive Systems (DIS '16). ACM, New York, NY, USA, 146-157. DOI: https://doi.org/10.1145/2901790.2901792*

95. *Amit Zoran and Joseph A. Paradiso. 2013. FreeD: a freehand digital sculpting tool. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13). ACM, New York, NY, USA, 2613-2616. http://doi.acm.org/10.1145/2470654.2481361*

96. *Amit Zoran. 2013. Hybrid basketry: interweaving digital practice within contemporary craft. In ACM SIGGRAPH 2013 Art Gallery (SIGGRAPH '13). ACM, New York, NY, USA, 324-331. http://doi.acm.org/10.1145/2503649.2503651*

97. *Amit Zoran, Roy Shilkrot, Suranga Nanyakkara, and Joseph Paradiso. 2014. The Hybrid Artisans: A Case Study in Smart Tools. ACM Trans. Comput.-Hum. Interact. 21, 3, Article 15 (June 2014), 29 pages. http://doi.acm.org/10.1145/2617570.*