OPTIMIZING MEDICAL MAKING: APPLICATIONS OF GENERATIVE DESIGN FOR FABRICATION IN HEALTHCARE SETTINGS

MEGAN HOFMANN
meganh@cs.cmu.edu

Human-Computer Interaction Institute
School of Computer Science
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213

July 11th, 2022

CMU-HCII-22-103

Thesis Committee:
Scott E. Hudson (Co-Chair), Carnegie Mellon University
Jennifer Mankoff (Co-Chair), University of Washington, Seattle
Patrick Carrington, Carnegie Mellon University
Jodi Forlizzi, Carnegie Mellon University
Björn Hartmann, University of California, Berkeley

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Human-Computer Interaction

Megan Hofmann: *Optimizing Medical Making:*
Applications of Generative Design for Fabrication in Healthcare Settings, © July 11th, 2022

FUNDING SUPPORT:
NSF: IIS-1718651, IIS-1907337, 1836813, 2031801, BPC-A #1246649, DGE1745016
National Institute on Disability, Independent Living and Rehabilitation Research: 90DPGE0003-01
DOD: Contract No. FA8721-05-C-0003 (or FA8702-15-D-0002)
Google Faculty Research Award, Autodesk,
CRA-W, Center for Machine Learning and Health, Siebel Scholars

KEYWORDS:
Accessibility, Digital Fabrication, Disability, Health, Human-Computer Interaction, Machine Knitting, Personal Fabrication, Personalized Medicine, Textiles, 3D Printing

When access is a practice of love it is no longer simply about logistics and something you have to do,
but something you want to do.

— Mia Mingus

If nondisabled Americans are getting the benefits of living in a world that has been designed for them,
then perhaps it is only fair that they compensate people with disabilities
for the challenges that the world produces. — Christopher Buccafusco

To my late Father, Steven Hofmann, who fostered my potential.
To the late Professor Adele Howe, who fostered the potential in so many.
Both set me on this path.

ABSTRACT

Disabled people and clinicians have a long history of inventing novel assistive technologies for themselves and others. This practice of creating assistive technology at the point-of-care is necessary to assure access to assistive technology. Disabled people, like all people, should be supported in augmenting their environment to their needs. Furthermore, their clinicians, as primary providers of assistive and medical devices, are best situated to deliver customized and personalized care.

It would at first appear that increased access to consumer-grade digital fabrication technologies like 3D printers would increase the scale and scope of these efforts. Unfortunately, the design tools needed to access digital fabrication are designed to primarily serve engineers, professional designers, or hobbyists. Such tools only leverage technical expertise, not the rich domain-expertise of disabled people and clinicians. A disabled person is uniquely qualified to express their needs and a clinician's medical expertise is often critical for designing assistive and medical devices that are verifiably safe and effective. Unfortunately, there is no way to express that expertise within current design tools.

I present two relevant bodies of research. The first is three qualitative studies of *medical making*, the practice of designing assistive and medical devices in a clinical setting or for use in a clinical setting. The second is the development of two computer aided design frameworks which support the (1) representation of domain specific expertise and (2) utilization of domain knowledge in generative design.

I focus on medical making, as opposed to general assistive making, because it is a critical access point where many people receive assistive devices. Medical makers form an new community of designers and digital fabricators whose ranks included clinicians, other healthcare providers, and disabled people. My qualitative research reveals that medical makers can recognize medical and assistive design requirements but struggle to use standard design tools to actualize a design that meets those requirements. That is, they can prescribe a solution, but they often struggle to make it. To meet the full potential of medical making, domain specific design tools would need to enable medical makers to specify their requirements, automatically generate designs to meet those requirements, and utilize design patterns from existing devices to inform this generative process.

I contribute two frameworks which enable medical makers to specify their intentions and generate designs that meet their specifications. My first framework, PARTs, scaffolds the 3D modeling process like object-oriented programming but in an entirely geometric, non-programmatic environment. This helps designers to add functionality and documentation to their design, amplifying their expertise as designs are shared and reused. My second framework, OPTIMUM structures the creation of domain-specific generative design tools, particularly those needed for medical making, around the mapping of objectives (i.e., designer goals) to modifiers (i.e., ways of reaching those goals). While OPTIMUM is derived from my design requirements of clinical CAD tools, it is also a generic and extensible framework which can be applied to a wide range of domains. I have developed five generative design systems using OPTIMUM in two medical domains, one assistive domain, and two fabrication domains. I present the final two systems, Maptimizer and KnitGIST in depth to demonstrate the capabilities of OPTIMUM as a toolkit for building domain specific generative design tools.

This body of work presents early findings in the burgeoning field of medical making. I recommend future work in the following areas. First, design tools should facilitate collaboration between patients and clinicians. Introducing new fabrication technologies into clinical practice puts clinicians and patients on equal footing as novices using these technologies. This may facilitate future collaborations between patients and clinicians that reduces medical paternalism and supports precision patient care. Second, as more 3D printable medical devices are created and shared in online communities, the potential risks of point-of-care facilities incorrectly selecting or producing a device could introduce significant medical risks. Medical making serves as a test bed for innovative medical device regulation through machine learning systems that can detect, and potentially correct, device risks or failures. Finally, medical making, like all fabrication applications, is limited by the materials makers can readily work with. In particular, textiles and soft materials are difficult to adapt to existing design tools. Extending the tool chains that support designing with a variety of material properties is critical to the future of medical making.

PUBLICATIONS

Ideas, writing, and figures have appeared my prior publications:

[1] Erin Buehler, Stacy Branham, Abdullah Ali, Jeremy J. Chang, Megan Kelly Hofmann, Amy Hurst, and Shaun K. Kane. "Sharing is Caring: Assistive Technology Designs on Thingiverse." In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. ACM, 2015, pp. 525–534. ISBN: 978-1-4503-3145-6. DOI: 10.1145/2702123.2702525. URL: http://doi.acm.org/10.1145/2702123.2702525.

[2] Erin Buehler, Niara Comrie, Megan Hofmann, Samantha McDonald, and Amy Hurst. "Investigating the Implications of 3D Printing in Special Education." In: *ACM Trans. Access. Comput.* 8.3 (2016), 11:1–11:28. ISSN: 1936-7228. DOI: 10.1145/2870640. URL: http://doi.acm.org/10.1145/2870640.

[3] Erin Buehler, Amy Hurst, and Megan Hofmann. "Coming to Grips: 3D Printing for Accessibility." In: *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility*. ASSETS '14. ACM, 2014, pp. 291–292. ISBN: 978-1-4503-2720-6. DOI: 10.1145/2661334.2661345. URL: http://doi.acm.org/10.1145/2661334.2661345.

[4] Megan Kelly Hofmann. "Making Connections: Modular 3D Printing for Designing Assistive Attachments to Prosthetic Devices." In: *Proceedings of the 17th International ACM SIGACCESS Conference on Computers &; Accessibility*. ASSETS '15. Association for Computing Machinery, 2015, 353–354. ISBN: 9781450334006. DOI: 10.1145/2700648.2811323. URL: https://doi.org/10.1145/2700648.2811323.

[5] Megan Hofmann, Lea Albaugh, Ticha Sethapakadi, Jessica Hodgins, Scott E. Hudson, James McCann, and Jennifer Mankoff. "KnitPicking Textures: Programming and Modifying Complex Knitted Textures for Machine and Hand Knitting." In: *Proceedings of the 32Nd Annual ACM Symposium on User Interface Software and Technology*. UIST '19. ACM, 2019, pp. 5–16. ISBN: 978-1-4503-6816-2. DOI: 10.1145/3332165.3347886. URL: http://doi.acm.org/10.1145/3332165.3347886.

[6] Megan Hofmann, Julie Burke, Jon Pearlman, Goeran Fiedler, Andrea Hess, Jon Schull, Scott E. Hudson, and Jennifer Mankoff. "Clinical and Maker Perspectives on the Design of Assistive Technology with Rapid Prototyping Technologies." In: *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '16. ACM, 2016, pp. 251–256. ISBN: 978-1-4503-4124-0. DOI: 10.1145/2982142.2982181. URL: http://doi.acm.org/10.1145/2982142.2982181.

[7] Megan Hofmann, Gabriella Hann, Scott E. Hudson, and Jennifer Mankoff. "Greater Than the Sum of Its PARTs: Expressing and Reusing Design Intent in 3D Models." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. ACM, 2018, 301:1–301:12. ISBN: 978-1-4503-5620-6. DOI: 10.1145/3173574.3173875. URL: http://doi.acm.org/10.1145/3173574.3173875.

[8] Megan Hofmann, Jeffrey Harris, Scott E. Hudson, and Jennifer Mankoff. "Helping Hands: Requirements for a Prototyping Methodology for Upper-limb Prosthetics Users." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. ACM, 2016, pp. 1769–1780. ISBN: 978-1-4503-3362-7. DOI: 10.1145/2858036.2858340. URL: http://doi.acm.org/10.1145/2858036.2858340.

[9] Megan Hofmann, Devva Kasnitz, Jennifer Mankoff, and Cynthia L Bennett. "Living Disability Theory: Reflections on Access, Research, and Design." In: *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '20. New York, NY, USA: Association for Computing Machinery, 2020. ISBN: 9781450371032. DOI: 10.1145/3373625.3416996. URL: https://doi.org/10.1145/3373625.3416996.

[10] Megan Hofmann, Udaya Lakshmi, Kelly Mack, Rosa I. Arriaga, Scott E. Hudson, and Jennifer Mankoff. "Making a Medical Maker's Playbook: An Ethnographic Study of Safety-Critical Collective Design by Makers in Response to COVID-19." In: *Proc. ACM Hum.-Comput. Interact.* 6.CSCW1 (2022). DOI: 10.1145/3512948. URL: https://doi.org/10.1145/3512948.

[11] Megan Hofmann, Udaya Lakshmi, Kelly Mack, Scott E Hudson, Rosa I. Arriaga, and Jennifer Mankoff. "The Right to Help and the Right Help: Fostering and Regulating Collective Action in a Medical Making Reaction to COVID-19." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445707. URL: https://doi.org/10.1145/3411764.3445707.

[12] Megan Hofmann, Jennifer Mankoff, and Scott E. Hudson. "KnitGIST: A Programming Synthesis Toolkit for Generating Functional Machine-Knitting Textures." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST '20. Association for Computing Machinery, 2020, 1234–1247. ISBN: 9781450375146. DOI: 10.1145/3379337.3415590. URL: https://doi.org/10.1145/3379337.3415590.

[13] Megan Hofmann, Kristin Williams, Toni Kaplan, Stephanie Valencia, Gabriella Hann, Scott E. Hudson, Jennifer Mankoff, and Patrick Carrington. ""Occupational Therapy is Making": Clinical Rapid Prototyping and Digital Fabrication." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. ACM, 2019, 314:1–314:13. ISBN: 978-1-4503-5970-2. DOI: 10.1145/3290605.3300544. URL: http://doi.acm.org/10.1145/3290605.3300544.

[14] Megan Hofmann et al. "Maptimizer: Using Optimization to Tailor Tactile Maps to Users Needs." In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI '22. New York, NY, USA: Association for Computing Machinery, 2022. ISBN: 978-1-4503-9157-3/22/04. DOI: 10.1145/3491102.3517436. URL: https://doi.org/10.1145/3491102.3517436.

[15] Udaya Lakshmi, Megan Hofmann, Kelly Mack, Scott E. Hudson, Jennifer Mankoff, and Rosa I. Arriaga. "Medical Maker Response to COVID-19: Distributed Manufacturing Infrastructure for Stopgap Protective Equipment." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445395. URL: https://doi.org/10.1145/3411764.3445395.

[16] Udaya Lakshmi, Megan Hofmann, Stephanie Valencia, Lauren Wilcox, Jennifer Mankoff, and Rosa I. Arriaga. ""Point-of-Care Manufacturing": Maker Perspectives on Digital Fabrication in Medical Practice." In: *Proc. ACM Hum.-Comput. Interact.* 3.CSCW (2019). DOI: 10.1145/3359193. URL: https://doi.org/10.1145/3359193.

There sat that beautiful big machine
whose sole job was to copy things and do addition.
Why not make the computer do it?
That's why I sat down and wrote the first compiler.
It was very stupid.
What I did was watch myself put together a program
and make the computer do what I did.

— Grace Hopper

ACKNOWLEDGEMENTS

This work could not have been done without the support of advisors, colleagues, family and friends. Thank you to my advisors, Jennifer Mankoff and Scott Hudson, for years of support, advise, and pointers to fruitful ideas. Thank you my thesis committee for their advise and support: Patrick Carrington, Jodi Forlizzi, Björn Hartmann. Thank you to Jamie Ruiz, Sarah Morrison-Smith, and Amy Hurst for supporting my undergraduate research. Thank you to my collaborators on works presented here and on the research detours I've taken along the way: Lea Albaugh, Rosa Arriaga, Cynthia Bennett, Xian 'Anthony' Chen, Jeffrey Harris, Gabriella Hann, Jessica Hodgins, Toni Kaplan, Devva Kasnitz, Jeeeun Kim, Udaya Lakshmi, Kelly Mack, James McCann, Ticha Sethapakdi, Stephanie Valencia, Lauren Wilcox, and Kristin Williams. Without your work, none of this would be possible. To my undergraduate and masters advisees: Nayha Auradkar, Jessica Birchfield, Jerry Cao, James Gan, Katie Lum, Anisha Nilakantan, and Emily Warnock. For stumbling around the dark with me, and often helping me find the light switch. Thank to my lab mates, both at CMU and UW, for creating the environment necessary to do this important work: Karan Ahuja, Nikola Banovic, Aashaka Desai, Kirstin Early, Jesse Gonzalez, Taylor Gotfrid, Dhruv Jain, Venkatesh Potluri, Daniel Revier, Michael Rivera, Anne Ross, Sai Swaminathan, Orson Xu, and Han Zhang. Finally, thank you to my family (chosen and by birth): Jon Schiavo, Elizabeth Hofmann, Steven Hofmann, Noah Smith, David Gould, Brian Chou, and Elena Chapman. I would never have gotten this far without you: discussing my research even when bored to tears or completely lost, making me eat and sleep rather than writing one more word, and generally helping me keep things together when I'm lost in this work.

The work presented in this thesis was funded by: the National Science Foundation (NSF) (IIS-1718651, IIS-1907337, 2031801, 1836813, DGE1745016, BPC-A #1246649); the National Institute on Disability, Independent Living and Rehabilitation Research (90DPGE0003-01); DoD Contract No. FA8721-05-C-0003 (or FA8702-15-D-0002) with CMU for the operation of the Software Engineering Institute, a federally funded research and development center; and my fellowships from: the Center for Machine Learning and Health (2017), the Graduate Research Fellowship through the NSF (2017-2022), and the Siebel Scholars Fellowship from the Siebel Foundation (2020).

CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACRONYMS

BLV blind or have low-vision

CAD Computer Aided Design

CDC Center for Disease Control

CHI Conference on Human Factors in Computing Systems

FDA Food and Drug Administration

FGO Functional Geometry Object

GUI Graphic User Interface

HIPAA Health Insurance Privacy and Accountability Act

IFU Information for Use

NIH National Institute of Health

OPTIMUM *Optimization Programming Toolkit Integrating Metaheuristic User-driven Methods*

OT Occupational Therapist

OTs Occupational Therapists

PARTs *Parameterized Abstraction for Reusable Things*

PPE Personal Protective Equipment

SCF Symposium of Computation Fabrication

SOP Standard Operating Procedures

TOG Transactions of Graphics

UIST Symposium on User Interface Software and Technology

VHA Veteran's Health Administration

1 INTRODUCTION

1.1 A CASE FOR ASSISTIVE AND MEDICAL MAKING

Brought on by the rise of consumer-grade fabrication in the early 2000's, maker culture or the "Maker Revolution" appears to be "one that the *average* person will be able to participate in and reap the rewards of participating...the revolution is open to *almost* everyone" [95]. Broad access to digital fabrication technologies (e.g., design tools, 3D printers) extends the flexibility of digital media to the physical world. But, critically, many people are still left out of this revolution; many populations that have a long demonstrated history of using innovative crafting and maker practices to benefit themselves and change the world—disabled people [142] and clinicians [84]. Even as costs go down, more advanced tools become available, and Computer Aided Design (CAD) tools become *easier* for novices to use, clinicians and people with disabilities are largely left out of the revolution.[1]

But disabled people have made and led revolutions of their own. Disabled people have consistently driven innovations in assistive technologies. Kudlick frames her book on "*Disability History*" [142] around the invention of numerous iconic assistive technologies (e.g., prosthetics, wheelchairs, white canes) by disabled people because it reveals how their innovative and inventive practices derive from a common need to participate fully in society. Wheelchair users in the 1970s modified and made their own assistive technology and simultaneously started a civil rights movement that fought for disabled peoples' right to access. Polio survivors who, in 1958, catalogued and published the modifications they made to wheelchairs and other assistive technologies [267], submitted their works as critical evidence in senate hearings to clarify what accommodation standards should be law [35]. Through individual, innovative acts of making, disabled people set the legal standards that guide accommodation practices to this day. Disabled people continue to make and innovate. Alongside their disabled contemporaries, clinicians, especially nurses, have been crafting and innovating at the point of care and have shared their designs since at least the early 1900's. From army nurses publishing sewing patterns for Do-it-Yourself gauze face masks to prevent the spread of the Spanish Flu in 1918, [15] to clinicians designing 3D printable face shields to fight COVID-19 [60], clinicians have always quietly made in the service of care. But generating and sharing designs remains a challenge; Gomez and Young get to the core of the problem

> *"The lack of recognition for maker nurses has a serious down side: All too often, nurse made solutions do not spread beyond the unit where they are created—or, even worse, are never built, remaining someone's interesting idea. " (Gomez and Young, [84])*

This rich history inspires two questions; how do novel fabrication technologies support or hinder the efforts of disabled and clinician makers, and what technical innovations are needed to amplify this work?

1.2 TOOLS FOR GENERATIVE DESIGN

Digital fabrication presents new barriers to disabled and medical makers because it relies on digital, inaccessible, and time consuming practices, like 3D modeling, that have been designed primarily to serve professional engineers. However, *generative design* is an alternative view of digital design that empowers makers to express their needs and domain expertise while automating away the technical work of crafting geometry and fabrication instructions. Unlike direct modeling, generative design uses designers' requirements and specifications to output designs that best meet those requirements. Rather than directly

1 Throughout the thesis, I use "disabled person" and "person with disabilities" interchangeably. "Disabled person" denotes a person who has claimed a disabled identity, (e.g., a designer who makes because they are disabled). A "person with disabilities" uses person first language to express that they are more than disabled (e.g., a person who is blind and uses tactile maps to navigate). I am a "disabled person" because my disabled identity is inseparable from my life and work.

creating a design, designers iteratively adjust their specifications to collaboratively generate a design with an algorithm.

Generative design is knowledge based, rather than skill based. For example, a blind person can generate a tactile map uniquely tailored to their needs with out using inaccessible 3D modeling tools (Chapter 9). Similarly, an occupational therapist with extensive knowledge of hand-splinting techniques could prescribe the right 3D printed splint for a patient without modeling the complex geometry (Chapter 8). Generative design is an excellent tool for amplifying the expertise of users while automating the time consuming and inaccessible work of digital design.

However, generative design tools often support automation by being domain specific and this requires that their developers have in-depth experience in both optimized search methods and the domain the tool will support—a rare combination. What is needed is a generalized tool for building new generative design systems in new domains. This tool must amplify the expertise of its expert users (e.g., disabled people, clinicians) through adaptable domain-space representations, tailorable objective functions, and intuitive systems for defining domain-specific heuristics. Ideally, this is all done while minimizing programming requirements through modular software construction and a Graphic User Interface (GUI).

1.3 RESEARCH APPROACH AND CONTRIBUTIONS

This thesis is composed of both qualitative studies of digital fabrication in clinical settings and the implementation and evaluation of toolkits which will advance new forms of computer aided design.

My qualitative work investigates medical making in terms of the prototyping practices of clinicians and the wider institutional infrastructures that support and hinder medical making. In Chapter 3, I used research-through-design methods to co-create 3D printed assistive technology with occupational therapists and their patients. Over this four month study, I gained critical insights into the barriers clinicians face in-situ to digital fabrication. Following on this work, in Chapter 4, I broaden this perspective by interviewing eighteen *medical makers*, digital fabricators in clinical practice. I found that these makers struggle to manage and mitigate risk when digitally fabricating designs made by others. This limits their ability to adopt designs shared online. Chapter 5 expands on this work during the acute phase of the COVID-19 pandemic (March-June, 2020) through an ethnographic study of the Make4COVID community which designed, produced, and deployed over 100,000 pieces of PPE to clinical facilities across the state of Colorado. This work demonstrates the challenges of translating the efforts of "expert-amateur makers" [143] into safety conscious clinical settings. From these studies I derive a set of requirements for *clinical CAD tools*.

Based on these requirements, I developed CAD tools that apply these requirements to new fabrication domains. My technical contributions include: (1) an approach to 3D modeling interfaces that enables makers to specify their intentions as well as geometry, and (2) a framework that supports the implementation of domain-specific metaheuristic generative design tools. Chapter 7 describes the 3D modeling tool, *Parameterized Abstraction for Reusable Things* (PARTs), which enables designers to 3D model as if they were creating object oriented programs. They embed their design intentions directly into the model which can be used to support heuristic based design variation in a generative design system. In Chapter 8, I present the *Optimization Programming Toolkit Integrating Metaheuristic User-driven Methods* (OPTIMUM) which enables designers to separate the search task into objectives which measure how well a design meets their goals and modifiers which strategically modify the design. Designers map these objectives to modifiers with priority rankings to create new optimizers that are tailored to their design goals. OPTIMUM (Chapter 8) is a generalizable toolkit for implementing domain specific generative design tools. I provide full demonstrations of its generality by optimizing tactile maps for blind users in Chapter 9 and machine knitted objects in Chapter 10.

In total, this thesis uses insights from two disparate domains to advance each other. From disabled people and clinicians, we learn the ways that groups of domain experts without technical expertise or practices can express their domain knowledge in reusable manners to create novel technologies and devices, which I used to structure a framework for developing generative design tools in new domains. From the field of generative design, we learn how to structure a design process as a heuristic based search with predictable and verifiable outcomes. These methods combine to support disabled and clinical makers in their endeavours to create safe, useful, and effective assistive and medical devices.

2 PRIOR WORK: FABRICATION FOR ACCESS AND HEALTH

This thesis builds on qualitative studies of maker communities and technical research on optimization methods and generative design applications in digital fabrication. An understanding of how people with disabilities and clinicians use digital fabrication technologies to create assistive technology informs the goals and requirements for useful generative design tools in the impactful domain of assistive technology design. Based on the prior research on making for accessibility and health, I argue that domain specific generative design tools are a critical innovation needed to amplify the work of disabled people and clinicians who already make assistive technology and medical devices. However, while there are numerous examples of bespoke generative design tools for digital fabrication, there is little research that supports the rapid development of domain specific tools. Such a generalized framework for generative design will be required to meet the wide range of needs of people with disabilities and caregivers. In this chapter, I separately review the relevant literature on (1) making for accessibility and health, and (2) generative design tools for digital fabrication.

2.1 MAKING FOR ACCESS AND HEALTH

Disabled people are not just the initial benefactors of universally beneficial technologies [80]; often they invent them. Glinert and York made one of the first arguments for "designing for disability" in computing, later termed *Universal Design* [145], by demonstrating that technologies that were made more accessible had universal benefits. While early adoption of the universal design framing ignored the contributions of disabled inventors, the argument often holds that innovations in accessibility have widespread benefits. There is a key difference between *design for disability* and *universal design*. Design for disability starts from the needs of a disabled person or group and may have external benefits, whereas universal design seeks generalized solutions by applying them to the specific experiences and needs of disabled people and other marginalized groups.

More so than a universal design model, the benefits of making align with the ability based design [268] framework. This approach emphasizes the value of customization. In particular, it calls for technologies to adapt to the abilities of its users through automatic customization features. While this approach has inspired significant work in user-interface adaptation using optimization methods (e.g., [43, 74, 153, 182]), there is little cross-over into digital fabrication research. Ability-based design's core requirement is that the act of customization is minimized, preferring to make technology automatically accessible. Making, alternatively, is a purposeful practice that elevates the disabled person from being passive-users to active-designers.

Beyond accessibility research, researchers tend to frame *making* (i.e., digital fabrication, consumer fabrication, craft) as an emergent form of physical end-user tailoring [211] and democratized design [244]. For example, Rosner and Bean studied the practices of "IKEA hackers" through interviews [211]. These hackers share their personalized modification and reuse of the popular global furniture store's modular parts. Rosner and Bean argue that these sharing practices reflect a convergence of online digital practices and physical prototyping. While the venues are novel (e.g., online repositories like Instructables [119]), the personal modifications described follow similar practices to the 1950 makers in Toomeyville [267]. Based on interviews with 14 hackers and makers, Hartmann *et al* describe making as a form of "opportunistic design" where "more time and ingenuity go to selecting components and shaping the glue ware that interfaces them" [94]. The focus of makers is on solving problems with what is available, rather than designing ideal systems.

Studies of online maker communities reveal common social values amongst makers. Makers at large respond to what Bean and Rosner term the "maker brand" [22] which positions consumers as producers and change agents. Kuznetsov and Paulos's review of online making repositories succinctly describes the maker values of openness, communal learning, and creativity over profit [143]. Tanenbaum *et al* [244] provides a useful expansion on this by describing digital fabrication as “democratizing” and “hedonist technologies” that “*privilege the pleasures of production over the value of the product*” [164]. Tanenbaum *et al* argue that digital fabrication enables craft practices to value pleasure in the act of creation, as well as afford the consistency and quality needed for profitability, which then enables new groups of makers to participate in the creation process. Lindtner *et al's* [157] focuses on the capitalist benefits of making, moving beyond personal benefits. Their study of makers' entrepreneurial endeavours shows that digital fabrication is adopted by novelty focused communities who wish to profit from the technical solutions digital fabrication technologies give them access to. Focusing on the products of digital fabrication can lead to a technosolutionist mindset, however Fox *et al's* [73] study of feminist hacker-spaces shows that maker communities centered around community and inclusion emphasize individual achievement over technological outcomes. Lindtner *et al* [156] noted similar critiques of making's technosolutionist brand but also mark the lack of an alternative “equally aspirational vision”. I argue that focusing on digital fabrication's role in supporting disabled and clinical makers highlights an alternative vision, one where these technologies help provision access in novel ways, but also create barriers to long standing practices where disabled people make for themselves and others.

When studying the intersection of making and disability, prior work has a tendency to describe their relationship in the wrong direction. That is, digital fabrication alone does not enable people with disabilities to make, rather disabled people have always been makers and are critical adopters driving the demand for novel digital fabrication tools. Numerous researchers have explored the intersection of digital fabrication and accessibility. They show that digital fabrication can be a powerful tool to amplify the efforts of disabled makers, but many barriers remain. From this body of literature we can primarily conclude that increased access to different design and manufacturing tools will empower more people with disabilities [117]. Enabling a wider set of designers and fabricators (e.g., disabled people, clinicians, care-givers) enables more people to design assistive technology for themselves and others which can enable the practical goal of increasing customizability, as well as the larger socio-political goal of empowering disabled people to participate more fully in the design and delivery of technology. In the next sub sections, I review the two bodies of literature on: (1) who makes assistive technology and medical devices and why, and (2) how novel CAD systems can support their efforts.

2.1.1 *Disabled Makers*

Broadly, there have been two primary approaches to studying making by and for people with disabilities. Many researchers, recognizing the relative rarity of disabled people using digital fabrication tools, have used research-through-design methods to engage with people with disabilities and other stakeholders in the making process (e.g., [26, 37, 53, 102, 171]). Others have used a range of observation, interviewing, and ethnographic methods to engage with maker communities that included disabled people [25, 96, 117, 204], volunteers [101, 113, 134, 191, 197, 237, 238], and clinicians [79, 84, 101, 167].

In his review of related rehabilitation literature, Quintero [207] lays out four implications for the field; that inclusion of people with disabilities in design: (1) grounds analysis of their needs, (2) increases our broader understanding of the field of rehabilitation, (3) facilitates knowledge transfer between disabled people and rehabilitation specialists, and (4) captures their subjective experience. While the influence of the rehabilitation field is clear, his discussion largely aligns with the findings from the broader HCI and Design literature on participatory research. For instance, it is widely accepted that the core benefits of participatory design are a grounded and inclusive view of participating users [216] and how the systematic act of design itself can reveal insights into under-constrained problem spaces [280, 281].

One core theme of disability and digital fabrication research that Quintero's review [207] does not discuss is empowerment. Hurst and Tobias established this theme as the basis for this field of research by examining the practices of disabled assistive fabricators over three contextual-interview case-studies [117]. Hurst and Tobias focus on the relationship between assistive technology creation and adoption, building on established concerns that assistive technology tends to be abandoned due to cost and poor fit to user needs [200]. Hurst and Tobias [117] found the making process empowered users to adopt their assistive technologies by forming a closer connection to them.

Meissner *et al* similarly cover the subject of empowerment but provide a more narrow definition [175]. Based on their review of online designs of maker-made assistive technologies, Meissner *et al* argue that empowerment is an exchange of opportunity between one group, presumably non-disabled, and a disenfranchised group (e.g., people with disabilities). They then argue the solution is an education process where they teach people with disabilities to make for themselves [175] without modifying the common technologies used for making to be accessible. While both works agree on the core benefits of people with disabilities making assistive technology, Meissner *et al* misses the countless examples of disabled people making for themselves. Further, their methods do not consider the specific access barriers posed by novel fabrication technologies (e.g., 3D printer [228, 229]) or other practical challenges faced by disabled makers [27, 116]. This makes their education-centric recommendations difficult to build upon.

In another study, DeCouvreur and Gossens studied making assistive technology in parallel to Hurst and Tobias, but they explicitly include clinicians as stakeholders in the co-design process of creating assistive technology [54]. They view digital fabrication as a bridge between a universal design philosophy [145] and the design methods used in rehabilitation engineering [207]. Like Meissner *et al* [175] and Quintero [207], DeCouvreur and Gossens argue that digital fabrication's value derives from the exchange of knowledge, which in their work is the implicit knowledge given by the disabled person to a clinician through the act of co-making. They argue that this exchange enables clinicians to produce innovations in classes of assistive technology distributed through clinical practice.

Other researchers have explored empowerment through social accessibility [227]. Shinohara *et al* coined the term *Social Accessibility* to describe how social stigma, particularly about the use of assistive technologies, creates access barriers for people with disabilities [227]. Bennett *et al* adopted this framing in their study of fourteen prosthetic users who modify their prosthesis primarily for aesthetic purposes [25]. Specifically, their participants who used open-source 3D-printed hands criticized the low quality of the designs while praising the social impact they had. These participants saw efforts to promote accessibility by using novel technologies like 3D printers as a way of promoting disabled experiences and voices. Sometimes these benefits have a positive impact on how people with disabilities view their disability and assistive devices. Proftia *et al's* study of an online forum of cochlear implant customizations revealed that customization helped wearers establish personal relationships with their assistive devices [205]. In a later workshop study with disabled activists, Bennett *et al* examined how the act of making can be used to protest the assumption that disabled people are not equal partners in design [26]. The "biographical prototypes" Bennett *et al* describe reveal a common bias that positions disabled people solely as a disenfranchised group, rather than telling the histories of disabled innovation (e.g., [35, 122, 142]).

While these examples address social inaccessibility through the aesthetics of the resulting artifacts, other communities of disabled makers address it through co-creation with non-disabled co-makers. For example, Das *et al* conducted an eight month participant observation of a weaving workshop at an assisted living facility for people with visual impairments [53]. The woven textiles made by the visually impaired residents were sold to the public to cover the costs of the workshops. This exchange prompted conversations between the weavers and the non-disabled volunteers about the perceived value of their labor. Those discussions lead to an understanding of the disabled residents' lived experience and reduced the volunteer's personal biases. This work expands our understanding of the benefits of assistive technology to include mechanisms for correcting social stigma.

While most of the literature focuses on disabled makers creating assistive technology (e.g., [25, 54, 117, 175, 204, 207]), other researchers have focused on how disabled makers "make making accessible" [116]. In mixed-ability maker spaces (e.g., [13, 53]) makers interdependently [24] adapt the space to fit each others' abilities and needs. Accessible crafting practices call for novel uses of materials that align with the abilities of users. For example, the blind or have low-vision (BLV) makers in Giles *et al's* study of accessible textile production emphasized the values of textural over visual aesthetics [77]. In the context of making prosthetics, my early research with people with limb difference showed the benefits of crafting with common materials (e.g., Legos, pipe-cleaners, Styrofoam) that are both engaging and easy to augment in a co-design activity [102]. In their report on "making nonvisually", Bennett *et al* [27], conclude that accessible making requires the elevation of different modalities and active check ins on access throughout the process.

Beyond the design of assistive technologies, making fabrication accessible is a critical step in including disabled people in all aspects of technology design and scientific research. Mack *et al's* [163] study of the accessibility practices of researchers shows a common thread between maker activities and research, in both cases we must "*anticipate and adjust*". Access is a constant labor, not just the delivery of an assistive technology. Thus, the tools that help provision access (i.e., fabrication tools) are both part of creating access and a potential barrier to it.

2.1.2 *Making for Disability*

The research described so far focuses on the practices of disabled makers, however researchers have also explored the practices of other assistive technology makers: volunteers [191, 197, 238] and, to a lesser extent, clinicians [37, 171]. In their review of assistive technology design posted on Thingiverse, Buehler *et al* found that only 24% of the designers that responded to their survey (25.6% of the 273 unique designers) had a disability [36]. The actual portion of disabled to non-disabled designers may be lower since the primary personal uses of assistive technologies (e.g., medication management, physical or motor impairments, and learning disabilities) each individually made up less than 1% of the design types, implying many non-disabled designers are making different types of devices than these respondents. Buehler *et al* also found that assistive technology makers tended to have technical backgrounds (48%), as the skills needed to use 3D modeling tools are often inaccessible, so many disabled makers have limited experience with these technologies. Critically, designers said they had dual motivations to make for others (a form of care-giving) and to overcome a personal technical challenge. Buehler *et al's* review demonstrates that there are two groups of assistive technology makers: (1) disabled maker communities that tend to use crafting techniques to overcome personal stigma (e.g., [53, 205]) and (2) non-disabled digital fabricators who see assistive technology as a technical challenge.

Studies of volunteer assistive technology maker communities have demonstrated that makers tend to exclusively focus on technical challenges in making at the expense of clinical, safety, and social concerns. The most studied assistive technology maker community is e-NABLE, a community started in 2014 that makes and distributes prosthetic-like 3D printed devices to recipients with upper limb-difference, primarily children. Parry-Hill *et al* provide the most comprehensive study of the community with a survey of the e-NABLE Google+ forum (63 responses) and 14 follow-on interviews [197]. Like Buehler *et al* [36], they found that few e-NABLE volunteers had clinical expertise or a disability. By contrasting these responses with those of clinicians and other volunteers, Parry-Hill *et al* demonstrated that volunteer maker communities do not have the tools necessary to foster effective collaboration with recipients or to customize devices to recipients' particular needs. Their findings echo those from my workshop study which brought together clinicians and e-NABLE community leaders [101]. In that study, we found that clinicians had significant concerns about e-NABLE's ability to regulate volunteer practices and ensure safety and product quality while operating in a open-community structure. The clinicians in the workshop highlighted concerns about the inability for volunteers to follow up with recipients or to methodically share results of technical experiments. Okerlund and Wilson found similar results in a smaller interview study with four college students from a local e-NABLE club [191]. Instead of contrasting these volunteers' perspectives with clinicians, Okerlund and Wilson emphasize how HCI methodologies could support these volunteers in building social relationships with recipients. From their perspective, the key limitation of these volunteers was their disinterest in forming relationships or a community with their recipients; all of their work focused on resolving mechanical challenges. E-NABLE is unique from educational and clinical settings because most volunteers have no direct relation to recipients. Without that direct social connection, these makers tend to focus on technical challenges.

Stangl *et al* presented an alternative volunteer organization which fostered closer relationships between recipients, care-givers, and volunteers through tactile graphic making workshops where volunteers 3D printed picture books for blind children [238]. Unlike the studies of e-NABLE, they found that volunteers had little difficulty collaborating with children and parents (some were care-givers themselves), but instead struggled the most with technical challenges and produced incomplete designs. These volunteer communities reveal a clear opportunity space to support volunteer makers in overcoming technical hurdles while maintaining a clear connection to the recipients they aim to help.

Making assistive technologies is attractive because it is a way to enact enact care for one's self, loved ones, or one's community [250]. Vyas [257] applied an ethos-of-care lens to his study on a crafting community of primarily women makers to evaluate makers' relationships within the community; he reveals a communal prerogative for altruism. Similarly, studies of makers of assistive technology reveal how makers attribute value to the tangible impact of doing good in addition to the personal benefits of the device [36, 198]. These makers follow the same altruistic directive as those Vyas [257] observed.

How care is practiced is dependent on makers' relationship with recipients. When the maker is both maker and recipient, the disabled person making for themselves [25, 26, 204], the practice of making is a form of self care. Alternatively, makers may make for someone they are close to as an expression of care for the loved one (e.g., a husband crafting a cane for a disabled wife [104]). Finally, making is one of many tools in a practice of professional care, like supporting students [37] or assisting patients [107, 172].

In recent years, digital fabrication technologies are playing an increasing role in clinical practice (e.g., [71, 127]), but little research has examined the challenges clinical stakeholders (e.g., people with disabilities, clinicians) face when interfacing with these technologies. Recall that Quintero [207] and DeCouvreur and Gossens [54] both applied a rehabilitation engineering perspective to their study of making assistive technology. In a rehabilitation model, clinicians are central to the delivery of assistive technology, yet neither body of work examines whether clinicians are able to deliver digitally fabricated devices at the point of care. Other researchers, like Ventola [255], have focused on examining the medical benefits of digital fabrication but not the challenges clinicians face. Ventola published an early review of the medical applications for 3D printing technologies and discussed four common benefits which align with the prior literature on digital fabrication for accessibility (i.e., (1) increased customization, (2) decreased costs, (3) fabrication speeds, (4) democratization of design tools) [255]. However, they never question whether clinicians can really make use of these technologies without the direct assistance of engineers and researchers. Research on clinicians making with or, more often, for people with disabilities often focuses on developing novel 3D printing techniques for prosthesis fitting (e.g., [39, 63, 64, 98, 232, 283]).

Most prior work has studied how clinicians interact with existing digital fabrication technologies without situating that experience in real medical practice. McDonald *et al* had physical therapy students create 3D printed assistive technology in a classroom workshop [171]. They found that clinicians adapted standard classes of assistive technology rather than inventing novel technologies. Additionally, they revealed a tension between the potential benefits of digital fabrication and clinicians' concerns about its adoption in clinical practice. However, McDonald *et al* provide few insights into how this assistive technology design process will involve real client-clinician interactions. Slegers *et al* conducted an extensive interview and workshop study with occupational therapists, similar to McDonald *et al's* methods, where the clinicians were asked to create 3D models using a common modeling tool [231]. Sledgers *et al* showed that, while clinicians consistently appreciate the potential of digital fabrication, they would not expect to find time to learn the required technical skills or fit the lengthy design process into already packed schedules. Aflatoony and Lee conducted a similar workshop but included industrial designer stakeholders [7, 8]. They found that clinicians appreciated the professional support, but the inclusion of professional engineers and designers is not feasible in most clinics. It seems that current maker processes are incompatible with the healthcare setting.

There are a few exemplar studies that included assistive technology fabrication with clinicians and disabled people. Buehler *et al* co-designed customized grips with occupational therapists in a school for children with disabilities [37]. Buehler *et al* proposes easy to use tools that support assistive technology customization. These insights were used to develop a tool for 3D printing grips [38] which they deployed at the school, and studied. The focus of this work was on the creation of assistive technology in an educational setting and the occupational therapists focus on their role in facilitating education. Buehler *et al* found that clinicians were receptive to the domain specific tool because it reduced the design effort and modeling skills required of them. Due to the small size of the case study, it is unclear if these results would be reproducible in general clinical settings with different clients. There remains an opportunity to examine how clinicians respond to digital fabrication technologies in-situ.

2.1.3 *Medical Making*

Medical making is the application of digital fabrication and craft skills by healthcare stakeholders to alter or supplement medical practice and infrastructure. It overlaps with communities that build assistive devices [36, 101, 198], focus on particular diseases [261], and repair failed infrastructure in times of crisis [82, 146]. Medical making occurs at many scales, such as do-it-yourself [25, 103, 117], between care-givers and recipients [104, 172, 189, 231], within medical institutions [84], and within grassroots communities [66, 82, 261]).

While clinicians making assistive technology is relatively rare, the wider medical maker movement is fertile ground for understanding how clinicians can interface with digital fabrication. Like assistive technology making by disabled people, making at the point of clinical care by clinicians dates back at least to the early 20th century [68]. Advocates of medical making, like Young and Gomez, argue that these practices have been largely ignored by modern medical research because the primary makers, nurses, hold less institutional power in healthcare [84, 273]. They argue that nurses and other frontline healthcare workers are the first to encounter problems at the point of care. Because of this they must innovate to meet patient needs by working through ad-hoc making processes. Glasgow *et al* 's [79] literature review of documented innovations by nurses reveals that many nurses struggle to translate their in-situ innovations into a wider change of practice, either locally in a hospital or broadly in the medical field. Glasgow *et al* present researchers an opportunity to engage with medical makers to better understand their longstanding practices.

The critical value of safety in medical making is, perhaps, the value that most differentiates medical making from other domains of maker culture. Making in a medical context exposes the maker and end-user to real and significant risks; it is not an exaggeration in some cases to say that lives are on the line. However, studies of medical making outside these clinical institutions rarely discuss mechanisms to reduce risks. For example, Parry-Hill *et al's* [198] study of the e-NABLE community showed that clinicians' attempts to inform other makers about safety-driven practices were rarely adopted by the wider community. While we would expect some risk-mitigation efforts to exist, mitigation mechanisms only add to the known concerns that organizational labor often falls on marginalized members of the community (e.g., women organizers [73]), nurses and other marginalized staff [146]). Organizers who may already be invisible in maker communities must enforce strict quality and safety procedures which do not align with makers' novelty oriented practices [105]. Because of this, medical maker communities that wish to broaden participation outside of clinical institutions may struggle to uphold safety as a core value.

2.1.4 *Tools for Making Assistive Technology*

Researchers have leveraged the emergence of consumer-grade digital fabrication to invent new materials and interaction techniques for accessibility. The majority of these innovations use digital fabrication tools to create new forms of tactile graphics for people who are BLV in a range of domains (e.g., educational models [222], screen navigation [130, 277], and maps [86, 109]). Fewer researchers have examined the design tools themselves. Instead of asking what tools can make, we must ask how target users will approach making with these tools. Some research has narrowly explored design techniques, rather than creating novel types of designs. Following recommendations derived from studies of assistive technology makers (e.g., [8, 117, 238]), these novel tools use generative design to simplify the modeling process by leveraging common patterns in specific domains of assistive technology (e.g., [34, 38, 45, 91]).

Recent innovations in materials and interaction techniques have primarily focused on using 3D printers and laser cutters as bridges between digital information and the physical world. There are numerous standard practices for creating tactile graphics for BLV users (e.g., the picture book makers described by Stangl *et al* [238]). A whole body of research has been built around automating these practices with digital fabrication. Kane *et al* seeded this research space with the invention of "Touch-Plates", laser cut tactile guides that could be placed on phone or tablet applications as tactile guides to a static user-interface [130]. This work demonstrated how laser cutting and 3D printing cheaply produces assistive devices to support BLV users when navigating inaccessible touch screens. Shi *et al* have presented a comprehensive exploration of this through a series of papers detailing different ways of labeling tactile educational models (e.g., a

3D printed globe with labeled continents) [220–223]. For instance, their tool, "Tickers and Talkers"[221], leveraged the acoustic properties of 3D prints to create machine-recognizable sounds based on different clicked locations on a 3D model. Following on this work, they create a digital labeling tool for these models that can be used by sighted makers to label models for BLV students [223]. Zhang *et al* [277] and Götzelmann [86] resolve labeling challenges by directly connecting the 3D model to common computing through capacitive touch screens. Both Zhang *et al* and Götzelmann *et al* create labels in their 3D models with conductive filament. When the user grounds the filament by touching it and pressing it against the touchscreen this creates a capacitive signal that triggers labeled responses on an independent smartphone app. Zhang uses this to create tactile extensions to smart phones like a calculator [277], and Götzelmann uses this to label maps [86]. Each of these tools create assistive devices through a narrow set of properties of carefully tailored 3D models, however this presents a challenge to researchers trying to generalize the results. Concepts such as using acoustic [97, 222] and conductive [86, 247, 277] material properties have been extended by other researchers, but the frameworks behind these systems have not been generalized.

One common framework that a few researchers have explored is generative design, where makers input specifications about the design (i.e., what they need it to do or be) rather than specifying the design directly. These specifications are then used to automatically generate the design. The simplest approach for generative design is to create models based on formulaic templates that makers fill in. For example, Brown and Hurst generate tactile graphics that represent math equations (e.g., bar charts, line graphs) by having makers fill in the equation and some specifications about the line thickness of the graphic, then the system generates a tactile version of the plotted equation [34]. The key benefit to this approach is that it is independently accessible to the BLV user, excluding the inaccessible printing process. While the BLV user cannot effectively model with inaccessible generalized 3D modeling tools, they can easily specify the plot desired and their design preferences through a screen readable interface. Taylor *et al* use a similar approach to improve on Götzelmann *et al's* [86] tactile maps designs [247]. Taylor *et al* has the user input a specific geographic region and map preferences and the system generates a corresponding tactile map for 3D printing. Unfortunately, maps are more complex than math plots and Taylor *et al's* users sometimes struggled to narrow down the map to the set of information they needed while still getting an easy to read result. These templated generative tools remove the skill of modeling but still require the maker to make all design decisions (e.g., information density, material properties). So, while it removes requirements, this alone does not expedite the design process enough to fit into clinical or personal making-practices.

One approach to removing the design work from assistive technology making is to use advances in machine learning and optimization to generate models with less effort on the part of the makers. Guo *et al* used advanced computer vision and machine learning techniques to create 3D printed tactile labels for appliances (e.g., microwaves) [91]. Guo recognized that the primary challenge for BLV users was to specify the location of buttons which required overlays; if they could see the layout they would not need the overlay. Guo *et al* pushed this effort onto crowd-workers who could train a computer vision system to recognize button layouts and then generate the corresponding 3D model. In the separate domain of grip design, Chen *et al* defined a parameterized space of common assistive gripping models and related them to different gripping styles for people with mobility impairments [45]. The resulting tool enabled makers to specify general approaches to interacting with an object that needed to be customized. The careful mapping of gripping behaviors to grip models creates a search-space that could be optimized over to generate an assistive device. Expanding on Chen *et al's* approach raises two challenges. First, how can the domain knowledge required to generate these mappings be effectively captured in new design tools? Second, once a mapping is established, what generalizable method can be employed to explore the resulting search-space in a range of diverse domains?

2.2 GENERATIVE DESIGN AND OPTIMIZATION FOR FABRICATION

Traditional CAD tools support designers in representing their designs, but not in doing the work of design. Generative design is an approach that can bridge the "*gulf of execution*" [188] between designers' real world goals and the challenges of modeling geometry. As prior work on accessibility and making has demonstrated, the process of representing designs creates the most barriers to domain-experts (e.g., disabled people, clinicians) working to make assistive technology. In traditional tools, designers must directly manipulate the low-level properties of representations of their design in order to achieve the functional goals of their design. Example manipulations could be laying out a circuit [14], documenting fabrication instructions [56, 147], or, most commonly, creating 3D model geometries (e.g., [16, 18, 50, 61, 275]). These representation manipulation tasks are often tedious and inaccessible [27, 116]. As an alternative, generative design tools may minimize the work of creating individual iterations by automatically searching for a design that meets the designer's objectives.

Generative design, particularly as it applies to digital fabrication, comes from the intersection of architecture, mechanical engineering, and computer science fields. The most practical generative design tools employed by architects (e.g., Grasshopper [89]) enable designers to define a space of geometry using a programming script, rather than a single geometry in that space [9]. Beyond defining geometry, generative design is primarily used to search for optimal solutions to designer-defined problems [49], a process called optimization. In Krish's review of generative design methods applied to architecture, they argue that generative design "*is structured to stimulate the designer's creativity by guiding the designer through viable design spaces constrained by performance criteria.*" [141]. That is, the optimization algorithm defines the space of designs possible and works with the designer to search that space for a solution to the designer's goals.

Generally, generative design helps designers explore a complex design space [23]. The results help designers brainstorm what is possible, but the optimized designs often need to be manually refined. While this can be a powerful tool for professional designers, there are many narrower domains where optimizing standardized designs could provide valuable gains to domain experts. In particular, automatically tailoring physical objects to end-user needs has demonstrated significant benefits for people with disabilities [165, 269]. The challenge of iterating-on and refining a standardized design to meet a patient's needs can create significant barriers for clinicians and medical makers. In these cases, a tool that automatically refines designs to meet domain specific goals can provide significant benefits. Creating domain-specific optimizers requires expertise from both programmers and domain experts.

2.2.1 *Generative Design Methods for Fabrication*

Table 1: Optimization Methods used in research articles published at CHI, SCF, TOG, and UIST from 2016 and 2021.

Venue	Percentage of Total Papers at Venue							Total Number
	Bayesian	**Convex**	**FEA**	**Numerical**	**Heuristic**	**Metaheuristic**	**Stochastic**	**of Papers**
CHI	10.3	6.9	17.2	17.2	20.7	20.7	6.9	**29**
SCF	0	16.7	33.3	0	33.3	16.7	0	6
TOG	3.3	56.6	10	10	6.7	13.3	0	**11**
UIST	0	8.3	8.3	8.3	33.3	33.3	8.3	**12**
Total Number of Papers	**4**	**21**	**11**	**9**	**14**	**15**	**3**	77
Percentage of Total Papers	**5.2**	**27.3**	**14.3**	**11.7**	**18.2**	**19.5**	**3.9**	

Across the literature, researchers have explored a wide range of digital fabrication domains for generative design (e.g., fabricating "surface like objects" [44], balancing 3D models [18, 124, 203], improving model strength [239, 265], or generating deformable mechanisms [16, 30, 50, 174, 201, 276]). We analysed 210 research articles from Conference on Human Factors in Computing Systems (CHI) [42] (140), Transactions of Graphics (TOG) [1] (18), Symposium of Computation Fabrication (SCF) [214] (6), Symposium on User Interface Software and Technology (UIST) [251] (52), published between 2016 and 2021 that included author keywords related to digital fabrication (e.g., fabrication, 3D printing, laser cutting) and optimization (e.g., optimize, inverse design, generative design). We then narrowed our analysis to the 77 papers that contributed a domain specific optimization method. We primarily excluded papers that either: proposed a novel fabrication method but no optimization algorithm; were studying fabrication practice but did not contribute a system; or described a generalized toolkit related to fabrication or optimization. We then inductively categorized the broader categories of optimization used in this body of work: Bayesian methods (e.g., [139]), convex methods (e.g., [187]), Finite Element Analysis (e.g., [278]), numeric optimization (e.g., [264]), heuristic methods (e.g., [5]), metaheuristic methods (e.g., [52]), and stochastic methods (e.g., [56]). We summarize this categorization in Table 1. Its notable that, with the exception of convex methods at TOG, no category of optimization methods are dominant. This means that toolkits that support any one of these categories, or, more critically, combinations of these categories, could support numerous domains.

Numerical and direct analysis methods are highly desirable because they are typically more efficient and provide provable guarantees about the optimality of designs they produce. Further, they are deterministic; they produce the same results for each configuration. In these methods, the objective function and constraints are formulated in such a way that the optimal solution can be computed. The simplest examples are systems of linear equations which can be solved with linear or integer programming by finding the intersection of the objective function and the constraint functions. These methods are often applied to variants of common mechanical problems (e.g., linkages). For example, Jiang *et al* used integer programming to solve a system of equations connecting beams that create large load-bearing architectural structures [124]. Few optimization problems can be formulated to use these methods because the models (e.g., linear sums) rarely describe complex systems. Most generative design tools that utilize direct analysis methods generalize the domain representation to a system of linked joints which can be analysed using forward or inverse kinematics [16, 17, 152]. Skouras *et al* used sequential quadratic programming, similar to linear programming but over systems of quadratic equations, to generate inflatable structures [230]. Where applicable, these methods are extremely powerful, both accurate and quick, but systems that can cleanly be defined by systems of equations are rare and difficult to recognize.

Another widely used set of methods are finite element analysis and topology optimization. Finite element analysis is a numerical method for estimating the forces on connected elements of a mechanical system, while topology optimization is a related optimization method that propagates changes to these elements across a system, with the goal of minimizing the estimated property. These methods have found numerous applications in mechanical and material engineering applications, including generative design of deformable input devices [16] and the construction of mechanical meta-materials [85]. In many cases, topology optimization is synonymous with generative design [215], however it has significant limitations for generating ready-to-manufacture results. While these methods produce high quality results in search spaces with large degrees of freedom, the resulting designs are rarely well-suited to common manufacturing techniques. Instead they are often used to generate concepts [23]. Guirguis *et al* survey challenges and opportunities to resolve this discrepancy [90] by combining topology optimization with evolutionary black-box methods. They find that domain experts play a critical role in adapting these methods to real world applications.

The most common optimization methods for digital fabrication in this survey were convex methods, such as gradient descent. Gradient descent is a powerful method for quickly finding local minima in the search-space, particularly over the continuous domain of 3D meshes. Usually, the optimization iteratively makes small changes to a design, essentially moving in small steps through the search space towards better outcomes (i.e., a local-minima). In many domains, these local minima are reasonable, potentially optimal, solutions. Further, unlike direct analysis methods, which require a vary narrow set of optimization function structures, many objective functions are differentiable or the gradient can be estimated using the Lagrangian method. Unique domains where gradient descent has been used include: masonry [265], balancing 3D printed models [18, 203], generating complex mechanism and robotic characters [50, 76, 276], and creating reflective surfaces [57]. The key limitations of gradient descent are that the search-space must be continuous and differentiable, and that the local minima be likely acceptable solutions. When analysing physical forces, this is often true. However, aesthetic and discrete domains (e.g., knitting [154, 185], modular furniture construction [252]) cannot be represented by differentiable objective functions. Further, designers have limited control over which local-minima are selected, and this may negatively impact designers' agency.

When enumerative, direct analysis, and convex methods are not available, generative design tools tend to rely on heuristic methods, which, in many cases, have been found to effectively produce results. Heuristic optimizations use information about the design domain (i.e., rules of thumb) to guide the search. For instance, when balancing an object, an effective approach is to put more weight near the center of gravity [203]. Heuristic optimizations greedily apply these heuristics with each iteration; that is, they apply the heuristic to make a locally optimal choice, without considering how it impacts the search globally. Researchers have used heuristics to generate models that meet requirements of stability [203, 252], material-usage [61, 159], and strength [239]. These can be particularly effective methods when insights into a narrow domain reveal clear heuristic patterns (e.g., airplane construction [253]). Heuristic methods trade-off providing guarantees about optimality for understandablity; that is, while none of these researchers make claims that their heuristics will lead to local-minima, a domain expert can readily understand how the algorithms make the decisions they do, similar to how the domain expert would make decisions as they directly designed the model.

When all else fails and the domain space does not lend itself to the aforementioned methods, there are a few examples of effective application of stochastic methods, which randomly sample the domain space (e.g., [28, 140]). The simplest stochastic methods (e.g., Monte Carlo methods) sample the domain space based on expected prior probabilities defined by the domain. More advanced methods blindly walk through the domain space initially, while iteratively learning and developing a probabilistic model which can be used to later narrow the search (e.g., Baysian optimization [138], Latin complement sampling [28]).

Metaheuristic optimization methods combine heuristic and stochastic methods to create high-level, problem-agnostic strategies to guide a localized search process [242]. Metaheuristic methods seem to be effective in a wide variety of fabrication domains, (e.g., [10, 52, 56, 72, 75, 88, 106, 121, 151, 279]). For example, a common metaheuristic method is simulated annealing which, rather than strictly following the gradient to a local minima, probabilisticly moves with the gradient to more widely search the space for a global minima [144]. Zhu *et al* , for example, use simulated annealing to create "motion guided mechanical toys" [279]. With this method they produce models of mechanical toys that match the motion paths provided by the designer. Another common approach is to use a parallel tempering method which conducts multiple gradient descent searches, seeded by random locations in the initial search-space, to find a set of local minima and select the best result [62]. Desai *et al* use this method to generate assembly-aware 3D printed cases for custom electronics [56].

Metaheuristic methods are more general because they can take advantage of the structure of a design space (e.g., differentiability, domain-specific heuristics) while broadly searching a wide space. These methods trade speed and guarantees of optimality for flexibility. This makes it a good candidate for helping domain experts and programmers with limited knowledge of optimization solve optimization problems. Metaheuristic methods can be inefficient and have no guarantees that they will converge on a global maxima or quality solution. Despite this, our survey of digital fabrication and optimization literature shows that they can be successfully applied to a wide variety of domains. They work in discrete, non-convex, and continuous domains and often they resolve to acceptable solutions. While many metaheuristic algorithms may be sufficient in a domain, selecting a high quality method requires analysis of end-user's goals (e.g., broad exploration of search space, efficient convergence). Blum and Roli explain that metaheuristics either "intensify" or "diversify" the search process using domain-agnostic mechanisms [31]. For instance, Tabu Search [81], diversifies by avoiding previously discovered designs, while Ant Colony Optimization [58] intensifies a search process by revisiting designs. Simulated annealing transitions from a diversifying strategy to an intensifying strategy. Even with an in-depth understanding of the domain, it is often difficult to know which strategy is best. Easily trying out different methods gives programmers the opportunity to find efficient and effective solutions.

2.2.2 *Designing a Optimization Algorithm*

These generative design tools demonstrate the broad promise of optimization methods in digital fabrication. Optimization methods can be used in a wide range of domains and automate the difficult work of design, however, most generative design tools are bespoke stand-alone tools that require extensive understanding of optimization methods from their developers. Without generalized tools for creating generative design tools in new domains, it is unlikely that this promise will reach disabled people and clinicians in assistive technology domains.

First consider how interactive tools can help define a search-space. All of the tools referenced in the last section either used carefully crafted representations of a narrow domain (e.g., shape grammars [14], collections of discrete connections with adjoining rules [124, 152]) or highly general domains such as 3D meshes. Narrow grammars are powerful, but they are difficult to construct over broad domains. Alternatively, generalized domains force programmers to manipulate representations that provide little domain-specific information (e.g., 3D mesh representations of furniture). Instead of making development easier, it leads programmers to struggle with evaluating and modifying designs.

Through her thesis [217], Schulz presents an alternative approach which leverages repositories of parameterized models, learning probabilistic distributions of their properties, which then affords more nuanced metaheuristic optimization methods. First, Schulz *et al* collect parameterized models of furniture and ascribe them to furniture templates which constrain how the parameters can be modified while ensuring the result remains a valid piece of furniture [218]. Then, they use this repository to pre-compute deformations of these templates and compare them to a wide range of desirable objectives [219]. These pre-computations create a probabilistic representation of the design space which can be stochastically searched. The main limitation of Schulz's work is that it only considers geometric parameters, and by extension, only supports optimization over geometric properties. There are many other properties of digitally fabricated designs that should be encoded in this search-space (e.g., assemblability [56]).

Even when a generalizable search-space (e.g., meshes) is used, information about these models can be collected and used to create variations in the model. Yumer *et al* , for example, takes a domain specific set of 3D models (e.g., shoes, cars) and presents them to crowd workers who label them with different subjective properties (e.g., fashionable, sleek) [275]. Then, their system learns the key aspects of the meshes that vary over particular domains and defines these variations as deformation handles. A new user can request a new model in the domain with specific properties and vary the deformation handles to find a solution to their request. For example, the deformation handle for creating fashionable vs practical shoes would identify changes in the heel height and new shoes can be made more fashionable by raising the heel. Yumer *et al* present an interesting approach to creating variations, the iterative step in any optimization process [275].

Schulz *et al* presents a solution to creating generalized representations of fabricatable designs [218] and Yumer *et al* a way of creating heuristic variations more easily [275]. The remaining component of Krish's breakdown of optimization methods [141] is to make defining useful objective functions simpler. In physical domains, the objective function represents some concrete physical property and users simply need to specify their goal (e.g., how strong should a model be), but subjective domains, like image editing [225], do not have clear objectives. Shimizu *et al* manage this with mixed-initiative methods where a set of random instances in their search-space are generated and the designer is asked to rank them. This ranking is used to create a probabilistic estimate of the designer's subjective objective function. New iterations are made searching for a solution to this predicted objective. With these subsequent iterations, the designer provides feedback by ranking new results and this feedback changes the probabilistic estimate of the objective function. The process repeats until the designer is satisfied with a result. The advantage here is that designers know the solution is right when they see it, rather than being forced to estimate their objectives a-priori. Unfortunately, an ill-defined objective makes it difficult for designers or developers to encode useful heuristics which will more efficiently search the space.

2.2.3 *Toolkits to Support Optimization*

Designing the best optimization method for a domain requires a design process with careful consideration, as well as iterative testing of each component of the optimization method. Programming toolkits can provide flexible systems to support this design process, rather than requiring programmers to implement each trial method from scratch. To evaluate optimization toolkits, we can consider criteria described by Olsen [192]: "*flexibility*", the ability to rapidly make design changes; "*expressive leverage*", the ability to "*accomplish more by expressing less*"; "*expressive match*", how well the toolkit model matches the user's mental model of the problem space; and "*ease of combination*", the ability to combine simple primitives from the toolkit into a wide set of complex solutions.

Fogarty and Hudson present a toolkit for user interface optimization that highlights these values [72]. Krish describes optimization as consisting of a search space or domain, a way of generating variation in that domain, and a method for evaluating designs in that domain [141]. A toolkit for optimization, ideally, gives developers flexible ways to express these components in a way that matches their mental model of the domain. Fogarty and Hudson structure optimization problems in a similar way, but focus on defining modular components of the optimization process that can be factored out to create useful toolkits [72].

In their GADGET toolkit, they break down optimization into: (1) initializers which define the starting point in a search-space, (2) iterations which step through the space, and (3) and evaluations of outcomes (i.e., objective functions). While both Krish and Fogarty and Hudson's deconstructions of optimization algorithms are useful frameworks for deconstructing design tools, they do not sufficiently decompose optimization problems so that domain-experts could craft their own optimizers using pre-made components. Tools building on these concepts are still necessary for creating domain-specific generative design tools.

Despite the wide adoption of optimization methods for generative design in digital fabrication, few toolkits have been widely adopted that make it easier to build these tools. In fact, across the 77 optimization papers we analyzed, only 2 specifically cited an optimization toolkit they used to implement their system. While researchers may be under-reporting the toolkits they use, this could also imply that existing toolkits do not meet user's needs when developing these systems.

Optimization toolkits tend to provide a set of general optimization algorithms, rather than recombinable primitives that enable programmers to experiment with and create their own solutions. Consider toolkits that support metaheuristic optimization (e.g., [194–196]) or optimization more broadly (e.g., [166, 249]). To access their optimization algorithms, programmers must carefully represent their designs in a standardized format, and expressing objectives and constraints can be equally specialized. This means that there may be a mismatch between how the developer understands their optimization problem and how they must express it. Further, not all algorithms require the same data, objective, and constraint formats. This makes it difficult to experiment with a variety of methods.

Alternatively, we can look to toolkits in the space of convex optimization [40] and Bayesian optimization [83]. Burnell *et al's* GPKit [40] was developed through an ethnographic study of a variety of engineering domain experts who use convex optimization. The toolkit is built around a highly flexible iterative cycle used to develop a convex expression of a design space, which enables the domain experts to carefully match their understanding of the design problem to a format amenable to convex optimization methods. Similarly, Golovin *et al's* Google Vizier [83] helps programmers tune Bayesian optimization methods to solve new problems where characteristics of the domain space are unknown (e.g., cookie recipes, user interfaces). GPKit [40] and Vizier [83] leave an open space for toolkits that support the optimization methods that make up 42% of our survey: heuristic methods, stochastic methods, and metaheuristic methods. These methods build on domain-expertise in the form of heuristics and are flexible enough to cover domains that are non-convex, discrete, and/or poorly characterized. This makes them good candidates for domain-specific generative design tools that can then be applied to medical and assistive domains.

2.3 SUMMARY AND CONCLUSIONS

Disabled people and clinicians have long been "stealth innovators" [273] of assistive technologies and medical devices. The emergence of consumer-grade digital fabrication technologies presents an opportunity for more makers to join these efforts and for researchers to examine these practices in new ways. While there has been a recent surge on studying makers of assistive technology, there is little work that deploys these technologies in clinical settings where disabled people can co-design assistive technology with clinicians. Further, the research has narrowly focused on assistive technology design, but does not consider the vast numbers of clinician makers when designing design tools for clinical settings.

Where researchers have invented novel design tools for assistive technology design, generative design seems to be a promising approach. There are numerous examples of generative design for digital fabrication outside of the assistive technology domain, but these tools are individually implemented with no generalized framework to guide the development of new systems. This generalized framework will be required for efficiently creating new domain-specific design tools that can meet the vast array of accessibility barriers fabrication needs to address.

3 "OCCUPATIONAL THERAPY IS MAKING": A PRESCRIPTIVE MODEL OF MEDICAL MAKING

Research at the intersection of digital fabrication and assistive technology rarely includes clinicians. The consequences of this are not fully understood, but may impact safety, quality, and even funding availability [101, 114, 171]. Additionally, digital fabrication is only accessible to a small set of people who own these technologies; it excludes people who primarily obtain assistive technology through clinicians. People may prefer a clinical model because they do not identify as disabled, the assistive technology treats a medical condition that requires clinical expertise, or they do not have the necessary technical expertise to make assistive technology for themselves. For people who primarily access assistive technology through clinicians, digital fabrication can increase access to customized assistive technology as long as clinicians can use CAD tools. However, little is known about how this technology influences clinical practice, or the challenges clinicians might face in adopting these tools. This chapter contributes insights about how digital fabrication influences and is influenced by clinical practice.

The CAD tools used in clinical contexts must be usable by clinicians. In low-resource clinics, clinicians are often the only person capable of creating assistive technology, and even when clinicians do have access to fabrication and fabrication experts, it is critical that clinicians can directly interface with appropriate CAD tools that can leverage their expertise and allow them to fully participate in co-design.

This chapter presents two case studies of assistive technology design by Occupational Therapists (OTs). Over four months, in two clinics, we provided fabrication expertise while the OTs managed their clients' treatment. The first clinic was free to uninsured clients and we worked with two OTs to develop a 3D printed thumb splint. The second site was a Veteran's Health Administration (VHA) clinic where we worked with two OTs to develop a knife grip and wheelchair transfer board.

We found that the OTs embedded digital fabrication into their usual client-care process in ways which were at odds with the rapid-prototyping practices that are characteristic of maker-communities [36, 197]. The OTs based their designs on standard classes of assistive technology and clinical expertise, rather than inventing new assistive technology. If necessary, they begrudgingly iterated on prototypes, but preferred to do so only if it resulted in a design that can be reused across many clients. They preferred to quickly adapt and customize designs using adaptive materials rather than digital iterations.

There was a disparity between the two sites based on their resources. The limited resources at the free clinic encouraged the OTs to adopt a maker ethos; they were excited to use common materials in unusual ways to support their clients. Conversely, the VHA OTs had many unusual fabrication resources and wanted to push the limits of what could be created with those tools, even when this drove the design process out of their area of expertise.

Based on these findings, we argue that clinical CAD tools should meet three design goals:

1. **Amplified Design:** Clinicians see themselves not as assistive technology designers, but assistive technology prescribers, prescribing variations on existing devices like doctors prescribe medication. CAD tools can amplify clinical effort by storing and distributing common designs.

2. **Appropriate Design:** Tools should help select the appropriate prescription based on available resources. This helps clinicians calibrate their expectations to what is doable, and may broaden their perspectives.

3. **Adaptive Design:** Tools should support adaptive modifications of a design rather than prototyping.

3.1 METHODS

We worked with OTs to co-design their clients' assistive technology in order to understand the benefits and limitations of consumer-grade digital fabrication technologies in a clinical context. Our methods are informed by participant observation [125], co-design [54], and research-through-design [281, 282]. Like Buehler *et al* [37], We note that 3D printing as a clinical practice was, at the time, too rare to study, and tools specialized to clinical fabrication practices did not exist. Instead, we propose a preferable future [281] where clinicians and clients can co-design assistive technology. To examine this potential future, we intervened [20] in the clinical context; as researcher and digital fabrication experts, we served the clinicians as a proxy for clinical-CAD tools. We evaluate this interaction [241], studying the relationship between CAD tools and clinicians.

We worked directly with OTs to design digitally fabricated assistive technology for their clients over four months. In clinic visits we: (1) directly observed the clinics' day to day operations; (2) conducted semi-structured interviews with each clinician to understand their decision-making process and the clinics' operating contexts; and (3) consulted with the clinician teams to digitally fabricate solutions with the potential of addressing their clients' needs. We encouraged the OTs to perform as much of the design activities as they could and only used researchers to support fabrication. We supplied fabrication technologies (e.g., 3D modeling software, 3D printers) and materials (e.g., printer filaments, carbon fiber) that could be feasibly accessed at each site. These case studies were contemporaneous, so actions taken in one study may have impacted the design activities in the other.

We collected twelve hours of interview and design session audio data. Notes and memos were connected to these audio records with a smart pen. We also took photographs and created design artifacts. Using thematic analysis [179], two coders developed 268 bottom-up codes then a third coder synthesized 27 axial codes. We collectively reviewed the artifacts, researcher notes/memos, and axial codes to develop themes. Several of these themes are dependent on a particular site's resources and clients. For instance, the lack of resources forced OTs at the free clinic to adopt a maker-ethos, while the abundance of resources at the VHA clinic encouraged them to push the limits of digital fabrication. As a result, we describe themes in the context of the most relevant case study.

Table 2: Summary of participant pseudonyms and roles

Site	Client Pseudonym	Clinician Pseudonym	Clinician Role
Free Clinic	Ron	Julie Sara	Instructor Student
VHA	Jon	Lorelai Anna	Practitioner Resident

We recruited OTs through word of mouth and had no relationship to them prior to the study. Even though we did not require it, all of the OTs had prior 3D printing experience. None had used it with a client. First, we met Julie at a free clinic for uninsured clients. She instructs occupational therapy at the nearby university and mentors Sara, a student who volunteers at the clinic. Next, we met Lorelai. She is the head Occupational Therapist (OT) in a VHA clinic and supervises the residents there. Anna is the resident who works most closely with Jon. The clinicians selected which client to work with, ensuring that we did not fabricate clinically-inappropriate solutions.

3.1.1 *Site 1: The Free Clinic*

We observed and participated in six clinic days while designing Ron's splint. Julie and Sara did the majority of the design work (i.e., sketching, thermoforming) and we translated the sketches into 3D models and printed them. The free clinic uses limited resources to support uninsured clients. Julie wryly explains:

"It's silly...we don't have...a splinting tray, so I've just been using a coffee pot " (Julie)

Whenever possible, Julie and Sara avoid prescribing assistive technology, preferring to prescribe exercises for pain reduction. When Julie does prescribe assistive technology, she usually prescribes a simple over-the-counter hand/wrist splint. Julie keeps a personal supply of these splints to donate to clients who cannot afford them. Because the free clinic cannot purchase assistive technology for clients, Julie personally covers the expense. We were initially concerned that 3D printers were beyond the reach of this low-resource clinic, however, Julie classifies printers as a medical tool for the whole clinic to use. In comparison, assistive technology for one client does not benefit the clinic as a whole so its the client's responsibility to pay for it.

3.1.2 *Site 2: The VHA Clinic*

The VHA has many fabrication resources such as local experts at a neighboring university and a VHA rapid prototyping laboratory in another state. Lorelai joined the study to learn more about fabrication so she can create assistive technology locally before sending requests to these facilities. She was unaware of the capabilities of these facilities and relied on local experts to guide her through assistive technology design that used these resources. The main limitation of these resources is their slow turnaround time. Lorelai described sending a wheelchair component to a VHA facility and waiting four months to receive a solution. The slow turnaround time and lack of interfacing with the client meant that the device was no longer relevant when it arrived and the resources were wasted.

> *"If you can make it, we can probably make it too. I just send it off to researchers at [a local university] and they will build it for me. The only problem is time...it may take them months to turn it around. " (Lorelai)*

We visited the clinic four times but never met Jon due to VHA policies. Our second visit took place during one of Jon's appointments and we met with the OTs in a separate room. Lorelai asked Jon questions and shared artifacts with us, running between each room. She viewed researchers as a stand in for the engineers at the rapid prototyping facility, so we conducted more of the design and fabrication activities than we did at the free clinic. Lorelai provided design specifications and sketches while we determined the fabrication methods, produced 3D models, and fabricated each design.

3.2 Results

This section presents themes through study narratives that derived from the design and fabrication of the two key artifacts: Ron's thumb splint and Jon's knife-grip. We summarize the themes as follows:

The importance of clinical expertise: Clinical expertise played a role in the design of each artifact. In each case, what the OTs designed was determined by the clinicians' perspectives on traditional assistive technology and ergonomics. This theme was made clearest by Julie and Sara's evaluations of Ron's thumb pain over time, and their impact on the design.

Cross-client reuse: Julie and Sara focused on creating a splint pattern that would be reusable with other clients. This same reusability was not necessary at the VHA, which had the resources to create highly customized and unique designs. For the free clinic, reusability excused the costs associated with a particular client, Ron.

Maker-OT identity: Julie highlighted the relationship between maker culture and occupational therapy when she reflected on her practice and how it has change over the years. Because of the limited resources at the free clinic, both Julie and Sara were creative in how they made assistive technology, in a way that they associated with a maker identity. In contrast, the VHA clinic had abundant resources, so Lorelai viewed the work as engineering and rarely called out a maker ethos.

Prototyping as failure: The OTs strongly rejected rapid prototyping (i.e., iteration on low fidelity designs) because of the cost to clients. Lorelai had already iterated on the knife-grip before joining this study, and saw those iterations as failures due to their lasting effects on her client, destroying his preferred knives and requiring multiple clinic visits over months.

Figure 1: These images show the stages in the design of Ron's Splint. Top left: his original gray, over-the-counter splint. Top right: the first iteration of the splint that was brittle and snapped when thermoformed. Bottom left: Sara thermoforming the final splint in a coffee pot. Bottom right: the final design that Ron took home.

FOCUSING ON ADAPTATION: Adaptable materials were effective at both clinics. Adaptation brings the design process into client-clinician interactions, rather than a separate, costly prototyping process.

3.2.1 *Case 1: A Customized Thumb Splint*

At the free clinic, we worked with Ron, an African American man with chronic thumb pain. Julie asked for a 3D printed, thermoformable splint that fit Ron's hand precisely and blended in with his black work uniform and dark skin.

3.2.1.1 *The Importance of Clinical Expertise*

Julie and Sara brought in a clinical perspective that identified the causes of Ron's hand pain and applied standard solutions. Julie set the goals of the design process to balance between anatomical needs and Ron's social needs. Ron's large hands, swollen by his injury, had stretched and ripped the one-size-fits-all splint (Figure 1, top left). Sara explained that he had aggravated the condition:

> *"He keeps moving his thumb. The splint is so loose! He doesn't give it time to heal. " (Sara)*

Julie was concerned that Ron's swelling had not improved because he would not regularly wear the splint. She hoped that an aesthetically pleasing splint would increase his prescription adherence. When Sara completed the first splint (Figure 1, bottom-right), she noted that the splint was much less noticeable than the traditional alternative, remarking:

"I was excited because you had a black splint and I had black Velcro. It is usually white with black Velcro. " (Sara)

Julie and Sara used their clinical expertise to identify the source of Ron's pain. Julie patterned the splint to target the outer thumb joint restricting movement enough to reduce damage while giving him a range of motion that supported his work activities, making it easier for him to adopt this splint. The OTs followed up with Ron after he reported wearing the splint every day for two weeks. The swelling had visibly reduced and he reported his pain at a five, compared to the prior ten.

Overall, Julie and Sara applied their clinical expertise to identify the root cause of Ron's pain and the barriers that kept him from using an over the counter solution. Balancing between Ron's medical needs (an immobilized thumb), daily living needs (moderate mobility), and socio-cultural needs (an aesthetically minimal splint) improved his condition.

3.2.1.2 *Cross-Client Reusability*

We developed the final splint over two appointments—an unusually quick turnaround for rapid prototyping. Despite this, Julie viewed it as a failure to require Ron to attend multiple appointments without receiving treatment, stating that it was the novelty of the research that kept Ron coming back.

"He's only coming back because he thinks you're cool. He wouldn't come back for just me and Sara " (Julie)

However, Julie hoped to save time with future clients by reusing the 3D model. In the first design session, we asked Julie to describe what she expected of 3D printed splinting. She presented samples of 3D printed splinting patterns that she had borrowed from a colleague and instructional materials which demonstrated traditional splint patterning on fabric and Thermoplast[1]. She annotated a splint pattern with Ron's reference measurements and asked us to create a parameterized model of the splint that she could reuse like her colleague's experimental 3D printed patterns.

OTs shape splints by heating and forming Thermoplast patterns in hot water and cannot iterate on the design without creating a new pattern. The more times Thermoplast is heated and cooled the more deformed it becomes, making it difficult to use. In contrast, Sara found 3D printed PLA easier to repeatedly reshape. She complained about the difficulty of using Thermoplast and encouraged Julie to bring PLA printing into her splinting courses. She especially liked the fact that the pattern was reusable after thermoforming.

"I love this pattern...it lays flat in the water instead of bubbling up...So much easier than Thermoplast " (Sara)

Over the next two appointments, Sara molded the different patterns to fit Ron's hand (Figure 1, bottom left). Each pattern had a slightly different shape, but all had the same underlying problems that led to failure. The patterns were too long, digging into Ron's wrist, and were thin and brittle (Figure 1, top right). Sara cut-off the bottom edge with splinting scissors and noted the correct length for the next pattern. We made small adjustments to the 3D printed pattern and were able to fit Ron with an effective splint during the next session.

While prototyping the splint across appointments was not ideal, Julie felt that the outcome was valuable beyond Ron's care. She now knows the relevant measurements on the hand and will no longer need to consider factors such as the thickness of the pattern, or how long the splint should be relative to the wrist. All of the information learned in this process is generalizable. She expects to reuse the splint pattern with other clients without the prototyping process from the study.

"Once we have the right pattern and know how to measure it, we wouldn't have to test it over so many sessions. " (Julie)

1 Thermoplast is a thermoformable plastic for creating splints.

Reuse occurred at two levels: (1) the OTs reused standard splint patterns and (2) the thermoformed PLA supported in-situ rapid reuse of the pattern. Julie viewed iteration on the 3D modeled pattern over two sessions as a failure, but hoped that it would be worth it in the long term because she could reuse that model. Sara, who interacted more directly with the materials, was excited about the ability to make quick in-situ adjustments to a standard pattern rather than having to re-create unique, but highly similar, patterns for each client. PLA and 3D printing lend themselves to quick reproductions of parameterized splints, while Thermoplast is exclusively manual.

3.2.1.3 *The Maker-OT Identity*

Throughout the process, Julie and Sara prided themselves on their ability to make creative use of common materials. For example, they used a coffee pot as a splinting tray because the free clinic has no other sources of hot water. This excited both of them because they believed creating assistive technology with their limited resources expanded the capabilities of the free clinic. When finishing the splint, Sara applied adhesive padding, and Julie sewed on Velcro straps. While Julie sewed the straps, Sara commented that she wished she knew how to sew and that it would make her work easier; Julie responded,

"Occupational therapy is making...you should learn" (Julie)

Despite this espoused maker-ethos, Julie and Sara were uncomfortable with 3D modeling in the clinical context. They understood and could use the tools but it did not align with medical software requirements. Julie pointed out features of medical records software to highlight differences between it and the CAD tools. She was concerned that there was no support within the CAD software for medical record keeping, maintaining client privacy, or ensuring clinician accountability. To her, the software was not appropriate in a clinical setting:

"So if you give me the files, what do we do with them? I can't just print them, what if he needs them?" (Julie)

Julie's maker-identity is superseded by her clinical goals: to protect and heal her clients within the regulations of health care. In this quote she is concerned that, while the CAD tools produced an effective splint design, she did not see a means to follow Health Insurance Privacy and Accountability Act (HIPAA) regulations [262]. Legally, a client's treatments must be documented and the client must be able to access that documentation if they choose to switch providers or get a second opinion, but Julie saw no way to ensure this with traditional CAD tools. It seems that Julie and Sara are comfortable with the craftiness of being a maker, but when that maker ethos enters a clinical space it is subject to regulation. The maker identity invites digital fabrication into clinical practice but it does not override the cautious do-no-harm philosophy of clinicians.

To summarize, the design of a thumb-splint at the free clinic revealed a tension between clinical and maker design methods. This contributed to Julie's concerns about the high cost of design iterations even though the result could support other clients. Julie and Sara's *clinical expertise* helped to identify the biological and ergonomic requirements of the splint in a way we would not expect of non-experts. The tools and adaptive behaviors characteristic of a *maker-identity* were necessary for the production of an effective and reusable splinting pattern. This splint *model reusability* made the prototyping behaviors acceptable. In the next section, that reusability is not present and the *consequences of a rapid prototyping* are more apparent, as are the benefits of *adaptive materials and design*.

3.2.2 *Case 2: A Modular Chef's Knife Grip*

Jon, a veteran with a spinal cord injury, was the focus of the study at the VHA clinic. Jon loves cooking but he needs customized assistive technology. Prior to this study, Lorelai made two customized knife grips out of Orthoplast[2] (Figure 2, top-left). Jon had limited success with these because they did not ergonomically align with the knife.

2 Orthoplast is a thermoformable plastic similar to Thermoplast. It is very difficult to remove from affixed surfaces. This makes a food safe grip, but prevents subtractive modification.

Figure 2: These images show the design progression of Jon's modular knife grip: Lorelai's original Orthoplast grip (top left); the initial pass at a 3D printed grip (middle left); a final design with red flexible inserts (bottom left, right)

3.2.2.1 *Prototyping as Failure*

Lorelai had iterated on Jon's knife grip prior to the study, but these prototypes failed, costing Jon his time and favorite knife set. From Lorelai's perspective, iteration represents failure because clients have no use for failed prototypes. Lorelai prototyped two knife grips over a series of months. Her first iteration of the knife grip (Figure 2, top-left) had a basic grip pattern with a guard loop to keep the knife in Jon's hand. This grip fit Jon best, however too much Orthoplast was up against the base of the blade, which blocked him from pressing the blade down completely. Lorelai created a second prototype on one of Jon's other knives. This version did not block the blade but the grip did not fit Jon nearly as well. Essentially, because the Orthoplast permanently modified the knives, he had lost his favorite knife and had only received ineffective grips. When we began working with Lorelai, Jon had abandoned the knives entirely. Lorelai hoped that by applying this research to the problem she could make up for Jon's lost knives and time.

3.2.2.2 *Focusing on Adaptation*

In response the permanence of the Orthoplast grips, Lorelai emphasized adaptability. During an initial design session with Lorelai, she laid out three key goals for the 3D printed knife grip. First, it would be food safe and hand washable. Second, it would closely match the geometry of the original grip and guard but accommodate Jon's ergonomic concerns. Third, it would modularly adapt to more than one knife in a set rather than permanently modifying his knives. Emulating prior research on 3D printing grips [38, 102], we printed the outside of the grip with black PLA plastic (Figure 2, middle-left). Inside the outer shell we added an "uncertainty buffer" [133] of red flexible nylon material which allows Jon to replace the knife without re-printing. Finally, we coated the components in silicone to make the grip soft and food safe. Lorelai can use this silicon coating to add or cut-away padding layers to better adapt the grip to Jon's hand.

This design demonstrates the tension between rapid prototyping, a practice ubiquitous in maker culture, and adaptive design. Each component of the design emphasized adaptability: the outer shell fits onto new knives, the "uncertainty buffer" adapts to different handle shapes, and the silicone supports adaptation of the fit of the grip. This adaptability distinguished the design from Lorelai's prior prototypes. The *negative consequence of rapid prototyping* in clinical practice is that a prototype that is not safe or effective could hurt a client or, at a minimum, could discourage them from adopting the assistive technology. *Adaptive design*, rather than prototyping, better reflects the iterative structures we observed.

3.3 Discussion

One model of medical making is to provide fabrication professionals and resources to clinicians, as we did in this study. However, even with access to state-of-the-art fabrication facilities and dedicated experts, it is clear from the VHA case study that current design methods and tools do not fully support clinicians. Further, many clinics, such as the free clinic, do not have the means to hire expert fabricators but could still benefit from digital fabrication. CAD tools that would support fabrication in these limited resource environments would have a broad impact on low-income populations—populations that include many disabled people who do not have the resources to access fabrication on their own.

For the potential of medical making to be met, clinicians need specialized tools to support their practice. We found that rapid prototyping processes assumed by many CAD tools do not translate into a clinical context. Maker culture has a fail quickly and take risks attitude. Not only is this in conflict with a do-no-harm clinical mentality [101], it ignores the disparity of resources and operating structures present in clinical environments. Iteration in clinical practice occurs at a macro-level, through carefully regulated research and development of new technologies, and at a micro-level of adapting a user's devices. We found few opportunities to present clients with low-fidelity prototypes, even if that would produce a better design.

3.3.1 *Amplified Expertise: A Prescriptive Model of Assistive Technology*

The OTs in this study did not view assistive technology creation as design or engineering, but as customizing prescriptions of assistive technology. After all, we did not invent a new type of splint or grip; we merely customized existing designs. There were two reasons for the prescriptive approach. First, the OTs began the process using domain-expertise in assistive technology, rather than considering what they could invent; their goal was to match clients to the best available technology, knowing it would be safe and effective. Second, a prescriptive approach takes less time, requires fewer resources, and poses fewer risks than inventing a new technology. It is better to provide a safe, working solution quickly, than a novel solution too late and after harm has been done. Finally, it also amplifies the utility of each design by ensuring its reuse across many clients.

3.3.2 *Appropriate Design: Resources Across Clinics*

Fabricating assistive technology in occupational therapy highlighted a relationship between design processes and available resources. While these considerations occurred at both sites, the disparities between the free clinic and the VHA revealed the limitations for low-resource clinics. The presence of a *maker-OT identity* at the free clinic revealed a willingness to bridge clinical and maker practices to make better use of limited resources. In contrast, the VHA OTs desire to *push the limits of consumer-grade digital fabrication* revealed challenges for clinicians to determine what can be done with consumer-grade fabrication and how to make best use of resources. Because of the low-income status of the clients at the free clinic, Julie based her design decisions on materials she knew were available and inexpensive to the clinic (i.e., PLA replacing Thermoplast, a coffee pot replacing a splinting tray). It was easy for the OTs to work with these materials which may have contributed to the success of Ron's splint. In comparison, Lorelai was relatively unconcerned about material cost or the fabrication process. Her challenge was knowing how to effectively use her fabrication resources. Lorelai and Anna had a general sense that the VHA could fabricate complex designs, but had no knowledge of what the fabrication process involved. This made it more difficult for Lorelai and Anna to reason about their designs.

3.3.3 *Adaptable Design: Iteration in Clinical Practice*

The common belief in HCI is that iteration is the core of design, necessary to "getting the design right" [41]. Even considering time and material costs, iteration still produces the best solutions [59]. As a result, CAD tools are nearly synonymous with rapid prototyping [19]. Despite this, the OTs rejected prototyping; they had one shot. When clinicians present clients with assistive technology, it does not need to be perfect, but it must be verifiably safe and useful. It is possible that prototype iterations would work in certain cases—especially when there is no urgent need for the assistive technology—but clinicians must trust the quality and safety of any prototype they deliver. Such high-quality prototypes seem antithetical to a rapid prototyping process; rapid prototyping, by definition, does not produce results on the first try.

3.4 DESIGN RECOMMENDATIONS

A culture of design and fabrication does exist in clinical practice, but it does not follow the rapid prototyping model found in non-expert, enthusiast maker communities [36, 197]. This requires a new set of clinical CAD tools that support macro and micro iterations on assistive technology design. We recommend that a clinical CAD tool (1) *prescribe clinical expertise* from medical research and practice through a *prescriptive library of assistive technology* that (2) is filterable by *appropriate design* characteristics based on *available resources* and (3) uses *adaptable tools and materials* to support modification in client-clinician interactions.

AMPLIFYING EXPERTISE WITH A PRESCRIPTIVE LIBRARY Clinical CAD tools should include assistive technology libraries that help clinicians to leverage their expertise and build on existing research in the medical community. Ideally, clinicians could search the library based on diagnoses. These libraries should update based on new research and contributions from multiple communities and stake holders (i.e., volunteer organizations, biomedical engineering groups, academic and industrial research). Clinicians could adjust models with medical information and client data using a generative design process.

APPROPRIATE DESIGN: MAKE RESOURCES SALIENT Clinicians must know what is feasible, enabling them to maximize their use of resources without overreaching and putting the patient at risk. The clinical CAD tool must make resources salient, such as the materials that are available and the insurance costs. Resource awareness must scale to resource-diverse environments and present context appropriate solutions.

ADAPTABLE TOOLS AND MATERIALS Clinical CAD tools should emphasize adaptation; starting with a prescribed model, the clinician can tweak the design using adaptable materials. Ideally, adaptation is physical not digital. When adaptation is a digital process, it must be quick and understandable by clinical stakeholders.

3.5 CONCLUSION

Bringing digital fabrication to occupational therapy demonstrates similar benefits to past research and increases access for people with disabilities. However, we found the OTs negatively regarded rapid prototyping. Instead, clinicians emphasized minimizing iterations at the cost of innovation. Their primary goal is to quickly deliver something that helps the client, and we must build clinical CAD tools around this constraint. Instead, clinical CAD tools must leverage the iterative cycles done by the broader medical research community and support adaptive design. The next two chapters expand on this work by examining global practices of medical making before and during the COVID-19 pandemic. In the place of digital fabrication research that focuses on enabling generalized rapid prototyping, the following chapters describe systems that enable expertise amplification as the first step towards effective clinical CAD tools.

4 POINT OF CARE MANUFACTURING: NEGOTIATING RISK IN MEDICAL MAKING

Moving beyond occupational therapy clinics, this chapter broadens our understanding of medical professionals who use digital fabrication in their practice. These established medical makers have adapted their local infrastructure to support and utilize digital fabrication at the point of care. In particular, we focus on the risks to patients, as well as the risk management strategies medical makers adopted prior to the COVID-19 pandemic. Designing alongside the OTs provided a framework for how clinical CAD tools could support the individual design practices of clinicians, while this study reveals the complex network of clinical stakeholders and how clinical CAD tools should be situated in order to minimize risk to patients, as well as reducing undue legal and ethical burdens.

Healthcare is an ecosystem with many stakeholders. Patients encounter doctors, nurses, and medical assistants at the point of care. Such clinical staff are acutely aware of the open problems in delivery of patient care [21, 68, 84]. In recent years, there have been a growing number of cases where these medical practitioners use digital fabrication to problem solve and meet medical needs. This problem-solving often takes place within hospitals and other medical institutions [167, 184].

When clinical staff create artifacts for patients, they explicitly commit to "do no harm" at every stage of the prototyping process [101]. Despite some Food and Drug Administration (FDA) policies to regulate medical device manufacture, this practice, for the most part, is completely without oversight [235]. It is no exaggeration to say that the quality of these artifacts can pose a threat to life and limb. To uphold product quality, systems in material practices for making—code, design schematics, and manufacturing—need careful design. However, there are gaps in our understanding of the resources medical makers can use to uphold product quality. Bridging these gaps is critical to help medical makers to work in a safe, accountable, and reliable manner.

To better understand the implications of medical making, we interviewed 18 healthcare stakeholders (e.g., clinicians, administrators, engineers, and medical researchers) across the U.S. and Canada who use digital fabrication in their clinical practice. We discussed their fabrication experiences, the process they follow, interactions with other stakeholders, and how making impacts their practice. Based on these interviews, we characterize medical stakeholders and their activity at sites of medical practice. We define *medical making* as using digital fabrication to modify medical processes and practices. Clinical Stakeholders who participate in the prototyping process of such artifacts are *medical makers*. The core contribution of this chapter is an in-depth analysis of the medical making ecosystem that predated COVID-19.

4.1 METHODS

The goal of this study is to better understand how medical makers leverage and create infrastructure to support and define their practice. This is the first study of medical making as it is integrated into point of care infrastructure by healthcare professionals, specifically those that work directly at the intersection of maker and healthcare infrastructure.

We conducted 18 semi-structured interviews with medical makers who actively used digital fabrication in their clinical practice. First, we analyzed publicly available information to recruit healthcare professionals who were advocates of the maker health movement. We interviewed different healthcare stakeholders: clinicians, administrators, engineers, and medical researchers. Between January 2018 and February 2019, we collected information about interviewees' maker technology experiences, their role in the making process, and their perceptions of how maker culture and fabrication affect healthcare. We gathered additional public data (e.g., news articles, blog posts, and social media data) to inform our understanding of stakeholder roles and ecosystems. We organized the data from interviews and public sources into inductive themes.

4.1.1 *Participants*

Table 3: Participants demographic data

ID	Profession	Specialty	Environment	Location	Gender	Patient Access
A1	Administrator	Neurology	Children's Hospital	USA	Male	No
A2	Administrator	Emergency Medicine	International	Canada	Female	No
A3	Administrator	Education Technology	University	USA	Male	No
C1	Clinician	Emergency Medicine	Academic Hospital	USA	Male	Yes
C2	Clinician	Endocrinology	Academic Hospital	USA	Female	Yes
C3	Clinician	Neurology	Academic Hospital	USA	Male	No
C4	Clinician	Cardiology	Children's Hospital	USA	Female	No
C5	Clinician	Audiologist	International	Canada	Female	No
C6	Clinician	Prosthetist	Private Practice	USA	Female	Yes
C7	Clinician	Public Health	Maker-space	Canada	Female	Yes
C8	Clinician	Occupational therapist	VHA	USA	Female	Yes
C9	Clinician	Occupational therapist	VHA	USA	Female	Yes
C10	Clinician	Radiology	VHA	USA	Female	Yes
E1	Engineer	Radiology	Children's Hospital	USA	Male	No
E2	Engineer	Rehabilitation	VHA	USA	Male	Yes
E3	Engineer	Rehabilitation	VHA	USA	Male	Yes
R1	Researcher	Public Policy	Government	USA	Female	No
R2	Researcher	Prosthetics	International	USA	Female	Yes

We sought out medical makers by reviewing news articles and social media, as well as attending makerfairs and fabrication conferences. While we sought both critical and positive perspectives on medical making, we only identified participants with positive perspectives. Because medical making is relatively rare and novel, there is a technosolutionist bias in the medical community towards digital fabrication [156]. Healthcare professionals tend to be either aware, positive, and active in making artifacts or neutral, unaware, and inactive. Before this recruitment phase researchers had no relationship to the participants with the exception of E3, who I had met at a 2016 workshop on 3D printing prosthetics [101].

Participants met three criteria: (1) they must be a practicing healthcare professional (clinician, administrator, engineer, or researcher) who are involved in patient care; (2) they must be based in the US or Canada, and be subject to the respective regulatory agencies; (3) they must use digital fabrication (e.g., 3D printing, programmable electronics, laser-cutting) to create physical objects for clinical practice.

In total, we interviewed 23 participants and excluded five candidates based on our criteria resulting in 18. Two candidates were excluded because they are researchers who study medical making but do not interact with patients. One candidate was excluded because he did not work in the U.S. or Canada. Two others were excluded because they do not practice medical making—their makerspace supports STEM education. The fifth candidate was excluded because she does not make physical objects in her work with clinicians. The remaining participants are described in Table 3.

In the rest of this chapter, we refer to the participants who are administrators, engineers, and researchers as facilitators. Facilitators are not clinicians, but perform integral roles in ongoing medical making activity. Figure 3 outlines the network of memberships available to medical makers.

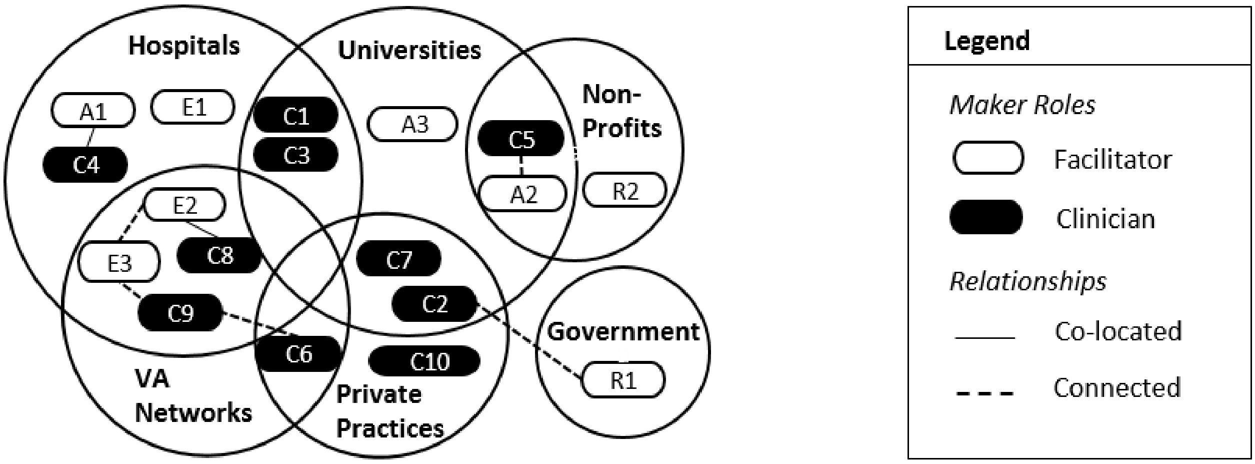

Figure 3: Study participants and their network of relationships with institutions and with each other. Proximity of shapes does not imply association.

4.1.2 *Data Collection*

The interviews elicited participants' thoughts on maker culture in healthcare as well as personal stories and experiences that could revealed their tacit beliefs about medical making. The interviews explored participants' experiences and perspectives on:

1. Their first and most salient making experiences in healthcare.
2. Their opinion of the role of fabrication in clinical practice.
3. Details about their maker space and fabrication technologies.
4. Their community medical makers and collaborators.
5. Their view of direction and scope of medical making.

With one exception, we interviewed participants over the phone or through their preferred video conferencing service. C10 was locally available for the interview. Interviews were audio recorded, with the exception of A1 who did not consent to audio recording. These interviews were between 32 and 83 minutes. Interviews were collected between May 2018 and February 2019. After each interview, the interviewer wrote a memo describing their initial thoughts. In addition to the interviews, we collected publicly available information that described participants' profiles and experiences with medical making. Researchers discussed the interviews in weekly meetings. Notes and memos led to a growing list of themes and new research questions which were the basis for a deductive thematic analysis of the interviews.

4.1.3 *Data Analysis*

We conducted an inductive thematic coding analysis on tangential data to the transcripts (e.g., public facing data that was collected on the participants, interviewer memos). Next, two researchers independently inductively coded the data to develop eighteen axial codes. These were deductively applied to the interviews. Our analysis showed strong inter-coder agreement (Cohen's Kappa coefficient (κ) = .777). Disagreements were discussed during the writing and synthesis process.

4.2 RESULTS

Our analysis revealed salient themes regarding the risks of medical making, stakeholders in medical making, and medical making infrastructure.

4.2.1 *Managing Risks to Patient Safety and Regulatory Gaps in Medical Making*

While all our participants were optimistic about the role of making in healthcare, several also expressed reservations about the risks of medical making. Medical makers evaluate risks of making devices based on regulatory guidance when available. However, in the absence of adequate guidance, each participant independently interpreted the risks. In general, participants who make devices for patient interaction were more concerned with liability arising from regulatory gaps, as clinicians are liable for the devices they create. On the policy front, regulatory bodies had formulated risk management policies [235] which have since been updated [3]. Meanwhile, practitioners developed their own strategies for managing accountability and protecting patients.

4.2.1.1 *Responsibility and Liability*

Medical makers are aware of risks to patient safety; the majority of clinician-participants raised these issues. C10 expressed her concerns about the consequences of on demand and small-scale medical making at points of care.

> *"The whole manufacturing side of it is very foreign to medicine...It's point of care manufacturing, bedside manufacturing. We are just not trained in that, nor do we think about all of the implications in terms of verification and validation." (C10)*

Clinicians bare most of the liability for devices. Participants who primarily make as a part of their private practice (R2, C2, C7) emphasized that clinicians may not fully understand the devices they deliver. C2 did not feel regulation of device quality was necessary when devices were made by and for individual patients. Her community fills gaps in mass-production of medical devices. However, she expressed a reticence to digitally fabricate without the support of engineers because she lacks technical knowledge.

> *"Perhaps you don't have the deeper understanding of the risks and benefits. Then someone else builds it and afterwards there's an adverse event...People who can benefit from [using 3D printing] are people who have the skill. How do you disseminate and democratize that?!" (C2)*

Facilitators are less liable for devices. For example, E3 creates devices that clinicians prescribe to patients, but clinicians are ultimately liable for his work.

> *"I don't have a license for rehabilitation engineering, right? That doesn't exist. On the liability side, [the clinician is] the lead therapist...I augment what they are doing." (E3)*

Regardless of position, medical makers took responsibility for medical making's risks. These medical makers go to great lengths to obtain licensing before distributing their designs. C7, C2, C5 and C8 each mentioned extensive time and energy devoted to license their designs. The devotion to following regulatory procedure differentiates medical makers from other volunteer groups or hobbyists who rarely, if ever, pursue regulatory review of their designs. C7 recognizes that it is a burden on medical makers to endure the licensing process, but that it is also a moral necessity.

> *"If we're looking to actually deploy it in global health settings, we work on getting FDA clearance of our devices. That's more resource-intensive but it's the right thing to do." (C7)*

4.2.1.2 *Risk Mitigation in Regulator Gaps*

Regulatory bodies have begun to consider medical making. R1 explained how amateur designs and maker creations are being explored at public health institutions. The National Institute of Health (NIH), FDA, and Center for Disease Control (CDC) each maintain maker spaces on their own premises to experiment with emerging maker technologies.

> *"[The FDA] are anticipating the moment when they are going to be asked to regulate all these things that makers are coming up with." (R1)*

Meanwhile, medical makers continue to uphold traditional healthcare ethics to the extent they can. Participants who operate in closed networks, such as the VHA (C8, C8, C10, E2, E3) or hospitals and universities (C1, C3, A1, R2, C6, A3), were more likely to support changes in regulations. The Canadian model of regulation allowed some participants (A2, C7, C5) to limit their scope of responsibility to device design. A2 defines their non-profit's making responsibilities within the confines of one-time production and dissemination of medical grade designs.

> *"Printing the pieces, distributing the pieces; that's not as much our project as much as testing them, validating them, publishing them, and then 'here you go world!'" (A2)*

Participants with patient access tended to seek out designs and tools with assurances of patient safety. Both R2 and C7 ensured a higher quality standard of production with standardized design repositories. C8 set out to find a tool that could be applied to her work in assistive technology and orthotics to measure outcomes and improve her custom designed devices with patient inputs.

> *"It's the only measurement that I found that assistive technology and orthotics kind of fit under...I was struggling finding one to assist me but that was the only one that I found." (C8)*

Other participants (C7, R2, C8) developed and shared guidelines to self-regulate quality. C8 collated a set of standards and guidelines from her VHA team, hospital partners, and the FDA.

> *"We are working on a 3D printer charter committee, where we are working to standardize how 3D printing is done in the VHA healthcare for various areas of healthcare." (C8)*

Similarly, C7 summarizes the ethos of medical practitioners as makers. She facilitates medical making in a space with procedures for quality checks and confidentiality agreements. In addition to valuing the products of medical making as intellectual property, she emphasizes her perspective on medical making, *"this is healthcare"*, a field separate from other hobbyist or general maker communities.

> *"I'm not trying to be secretive, but this is healthcare. I should not lose my medical license because of a maker's project. We [makers] have to be very vigilant about protecting [patient] privacy." (C7)*

Overall, medical makers were wary of the risks and their individual responsibilities when making. Collectively, participants developed a variety of resources and guidelines for risk mitigation. However, due to a lack of regulatory guidance, liability defaults to clinicians even when facilitators play a critical role in designing and adapting devices. This liability becomes even more complex when designs are openly shared, separating the maker from the practitioner. Because of this, many participants avoided adopting or sharing designs, especially when they needed to be adapted to a patient.

4.2.1.3 *Patient-Adapted Device Limitations*

Licensing design files can be useful in creating global resources for making. However, it may not fully address the challenges of assuring device quality when devices are adapted to specific patients. Participants had varied opinions on whether designs that needed to be adapted should be shared openly. Participants who are a part of non-clinical maker ecosystems (C1, A3), or who operate in non-traditional medical systems (e.g., non-profits) (A2, C5) championed openly sharing resources. In opposition, R2 and C7 would argue that even the best documented and tested designs cannot be adapted without patient access. In fact, R2 felt strongly enough to go where the patients were to pursue her experiments with 3D printing prosthetic devices in the Global South. She began medical making through a non-profit and decided to move away from a non-clinical volunteer community because it did not align with her focus on a process that prioritized patient safety.

> *"I sort of pivoted away...from the download things online and print them out anywhere, never seeing the patient sort of situation." (R2)*

Similarly, both E3 and E2 are making highly customized devices for patients with disabilities and do not distribute their designs. As E2 notes, this means that their work is not subject to FDA regulation. This leaves engineering labor in healthcare in a regulatory blind-spot where it is difficult to enforce best practices for patient safety.

"Assistive technology flies under the radar of the FDA." (E2)

Though we found no consensus among our participants on how the existing regulatory structures can accommodate making, participants acknowledged the need for quality in device production. Some participants (C1 and A1) argued medical devices could still be made within the existing regulatory structures. C1 shares his view that clinicians can make patient-centered interventions and need not be reticent due to concerns for patient safety.

"Everyone thinks you need to get FDA approval...people don't even [make] because it's way too complicated...We can't just make something and tell patients that its reasonably safe...There are still many more opportunities to create something that is not going to harm a patient." (C1)

4.2.2 *Leveraging Stakeholder Expertise from Medical Maker Networks*

Unlike medical liability, manufacturing expertise is distributed across medical makers. Medical making spaces include clinicians from many specialties, non-profit organizers, government officials, entrepreneurs, students, and hospital administrators. Irrespective of the site of healthcare delivery, clinicians and staff expressed an inclination to create solutions when supported with tools, skills, and other material resources. Participants shared their motivation to improve practice, deliver patient-centered care, and impact social good. They maintain memberships in several communities and advocate making practices in their institutions.

Medical professionals and facilitators shared a consistent motivation to improve their practice. A few clinicians (C1, C2, C3, C7, C8) and most facilitators (A2, A3, E3) mentioned publication goals. Clinicians with patient access (C2, C7, C5) highlighted making as a practical tool for innovating in routine care delivery. C7 described their low-cost prototype of a high-precision diagnostic device to replace a market option that was too expensive for wider adoption. Facilitators A2, A3, and E3 mentioned 3D printing as a process to create devices unavailable in certain markets. Similarly, C5 and C7 hinted at entrepreneurial innovation in their discussion of intellectual property rights. Further, a holistic intention towards social good guides their making practices. C2, C1, and R1 remarked on the state of public health in the US while A2, R2, and C7 offered a global perspective. C1 recounts his decision to adopt making as a medical professional and educator.

"I started thinking of the healthcare system as a whole and how broken it was and try to see how we might be able to fix it." (C1)

Clinicians act on their intentions by enlisting engineering expertise (e.g., C3, C8, and E2 collaborating within the VHA network). All facilitators in maker spaces (A2, C10, A3, A1, C8) referred to specific members with engineering responsibilities. Unfortunately, specialized medical engineering labor and technology is scarce—it is rarely available as a dedicated resource in medical settings. To bridge medical making needs, some participants (C2, C8, C6, C8, E2) use tele-health or remote consultations to share resources across locations. Engineers E2 and E3 leverage the VHA's resources to work with appropriate technology vendors for high-quality prints. Others in non-profits (A2, R2) and individual practitioners (C1, C2, C7) rely on academic partnerships to supplement engineering skill. A2 remotely co-ordinates project collaborations with engineers located in a low-resource setting, as well as several global partners. E2 shares his role in providing support to medical makers through tele-health in the VHA.

"So I do quite a bit of tele-health right now...And the reason for that is there are six rehab engineers within the VHA. There's probably only like three or four sites that are doing any sort of 3D printing clinically in the assistive technology area and not all of those sites have engineers." (E2)

Unlike many hobbyist makers, gaining technical expertise is not a prerogative for medical makers. Instead, medical practitioners (C1, C8, C8, C3) enlist technical colleagues to supplement their limited expertise. C8 describes how clinicians and prosthetists can collaborate and exchange expertise to create customized assistive devices.

"[Therapists] know that [3D Printing] is now a tool in their tool box that they can either hand over to [prosthetists] and we can come up with a solution." (C8)

Medical making is distinctly collaborative and involves teams with diverse expertise. Participants developed communities and collaborations between clinicians and engineers, and adopted technologies to facilitate these collaborations even when local expertise was limited. Rather than focusing on learning new disciplines, most participants preferred collaboration as a way of including expertise in their making activities.

4.2.3 *Facilitating Maker Operations in Medical Practice*

Medical makers are highly motivated individuals, yet they depend on access to technology and skill at the site of clinical practice. Co-located access within health institutions requires justification of setup and ongoing costs at the place of practice. The technical expertise required to adequately equip and repeatedly plan for medical and operational needs also poses ongoing challenges for facilitators. Our participants developed ways to organize making infrastructure around medical institutional practice or identified alternate means of making outside their institutions.

While several researchers highlight the benefits of setting up a makerspace in hospitals [167, 274], the high cost of hospital space makes it difficult for even well-funded participants to sustain maker spaces. E2 explains how setting up the space requires consideration of adequate room for ventilation and access for several stakeholders, such as a waiting room for patients. C7 was skeptical of co-location due to such cost barriers. In a similar vein, C10 described their solution to justify the allocation of space by sharing it with a research lab.

"Space is very expensive in a hospital. It's very hard to get space....Now our makerspace is in a lab. We share between research and clinical [sic], we joined forces." (C10)

The resources and equipment allocated to medical maker spaces are dependent on their institution's goals. Makers who prototype products (C1, C2, R1, C7, A3) and low-resource non-profits (A2, C5, R2) tended to use consumer-friendly machines. Other makers who worked in specialized clinics mentioned specific 3D printers and materials for life-like medical modeling (A1, C4, C3, C10) or assistive technology engineering (C6, C8, C8, E3, E2). C6 described the challenge of sourcing appropriate 3D printers for her prosthetic clinic because she requires specialized equipment that produces stronger and more precise models than consumer grade devices.

"We do not have any 3D printers now. About two years ago they were given to us by the cardiology department but it wasn't the right kind for prosthetic devices. We couldn't use them clinically." (C6)

For some participants, the choice of technology is tied to the operational costs of materials and making in low resource conditions. While C10 explains how she initially invested in machines at the VHA lab space without considering the cost of materials, others (R2, C5, A2) with fewer resources rely on constraints to guide their technology use. R2 explained the effect of electricity costs in the Global South. She limits design print times to under six hours by adapting the design into smaller modular parts. C5 similarly explains how 3D printing supplements their organization's limited access to medical devices.

"One of the problems with a place like ours [is that] you can't get the materials in there to make a [medical device]. 3D printing gives you the means to melt down plastic and build from like seemingly nothing." (C5)

Once technology is in place, several participants fund or sustain materials for making. Our participants employ different strategies, including institutional advocacy (C10, E3, A2, A3, C1) and grant applications (C8, A3, C3, A2). Others adopt practices often seen in maker culture to lower operational costs through the use of open-source software (C1, C5), crowd-sourcing skill (A2, C7), or crowdfunding to meet the cost of materials (A3).

> *"We have a crowd-sourcing project that we launched in the summer where people from the [university] and others outside can donate money specifically for funding our filaments and resin. That takes a bit of that burden off the library." (A3)*

In summary, the adoption of making within organized medical practice requires our participants to overcome several infrastructural barriers. Our participants mobilized resources for making devices at the point of care with a pragmatic attitude. C2 expresses an underlying derision in her response to a question about the challenges she had faced in maker projects for health.

> *"There's so many barriers to the craft of design inside the delivery system that, even if you are a provider and you know what the solution can be, you don't have any means of actually creating or supporting it." (C2)*

4.3 DISCUSSION

This chapter presents an analysis of medical making practices to distinguish the risks of applying digital fabrication in professional medical practice. These practices inform the values clinical CAD tools must support: safety and expertise-diverse community.

4.3.1 *Safety-Centric Making Practices*

Medical institutional norms supersede maker culture's norms in medical making. The ideal of "do-no-harm" is enforced top-down through regulatory structures that are, at present, unable to adapt to medical maker's practices. Medical devices are risky propositions for patients, even when they are sufficiently regulated [69]. Regulatory policies become inadequate and ultimately opaque when applied to making in medical practice. To compensate, medical makers adopt new ways to guard against risks.

Medical makers hold themselves to clinical standards in order to uphold patient safety. In effect, they are healthcare providers who apply the opportunistic design practices of makers [94, 211] to deliver healthcare. Such an extension of the physician's role alters the framing of making as a practice by "expert amateurs" [143]. Medical makers are healthcare experts, but making entails adding skills to their expertise to care for patients. Medical makers perceive making as a tool with the potential to make a significant impact at the point of care. However, like any other tool, it cannot rely solely on skill and attention of the user to ensure a higher quality outcome.

Medical makers in our study were motivated to produce verifiable, safe, and effective medical devices for patient care. Regulatory gaps leave clinicians liable for facilitator's actions. The medical institution and practitioner are legally responsible for medical device manufacture and distribution. In contrast, engineers play a critical role in ensuring product quality, which is not reflected in their legal liability. Such a skewed distribution of risk and responsibility exposes medical practitioners to malpractice.

Overall, medical makers take the additional responsibility to distribute universal designs. Medical makers adopt, despite uncertainty in regulatory guidelines, pragmatic approaches to mitigate risks. Medical makers attempt to adhere to regulatory policies by creating extensive documentation to safeguard design customization to the greatest extent possible. Unlike hobbyist makers that tend to provide minimal documentation [115], we may be able to rely on medical makers to provide extensive documentation of their designs. They ensure adequate documentation of the manufacturing process is available in text, code, video, and other media formats to invite feedback. Such processes are currently initiated and upheld by the medical makers involved in this project. There is an opportunity and demand to standardize these practices through collaborative platforms and policies.

4.3.2 *Building a Medical Maker Community*

Medical making requires clinical and technical experts to cooperate in the prototyping process. The collaborative structure of medical making is key to creating medical devices that meet the high standards of "doing-no-harm". Medical makers were quick to identify and proclaim their inadequacies, and to seek assistance from other medical experts who could resolve them. We found that consistent collaboration across institutions and individuals sustains a culture of innovation at the point of care, and that medical makers maintain a global and local network of collaborators.

While medical makers prefer locally available technology expertise, it is not always possible. In such cases, resources are made available remotely through institutional networks (e.g., VHA). Unfortunately for medical makers in smaller or less supportive intuitions, cross-institutional infrastructure is rarely available. This suggests that a tool to manage this process could improve remote collaboration and labor distribution across multiple medical institutions. Further, such a tool would have the benefit of an abundance of designs which could be evaluated automatically.

4.4 DESIGN RECOMMENDATIONS

These participants highlighted their strategies for risk mitigation and expertise collation. Some of these strategies rely on the online, open exchange of designs, similar to the practices of wider maker culture [143]. In this section, we propose two design recommendations that expand the requirements for clinical CAD tools: (1) support partially open distribution of designs which meet regulatory standards, and (2) develop a wider network of medical makers within and across institutional boundaries.

4.4.1 *Supporting Medical-Grade CAD Designs*

Thorough testing and documentation of a design is one of the most tedious, difficult, and costly parts of any engineering or maker effort. It is also absolutely essential to medical making. Even within hobbyist maker communities, design documentation is a critical factor in the distribution and reuse of designs [6]. Given the heightened risk of making in a clinical and patient centered environment [101], it is clear that medical making requires a clear and usable documenting procedure.

Thus, clinical CAD tools will require that medical makers can easily document their designs as they create them. This documentation process should be integrated with direct modeling and design creation and scaffolded to consider common regulatory concerns related to patient safety, design efficacy, and a reproducible manufacturing process. Clinical CAD tools should include documentation templates that inform makers about these broader concerns and create a common format for documenting these concerns. These templates should scaffold regulatory review and facilitate clinician adoption and adaptation.

4.4.2 *Distribution Networks within the Medical Practice*

Documentation of models can help in all contexts, but it is critical to support a wider distribution of medical designs across communities of practice. However, current open-source repositories for digital fabrication [36, 119, 243] do not provide sufficient infrastructure for sharing medical maker designs in safe and effective ways. In addition to the clinical CAD tools, medical makers require a *Medical Maker Repository*, such as the NIH 3D print exchange [118]. The repository must differentiate between prototypes and regulated designs so that clinicians can quickly adopt devices that are appropriate for their application. Unlike the NIH 3D print exchange, this repository must help contributors distinguish designs at different levels of fidelity and testing. It is notable that such a structure was adopted by the NIH for their review of COVID-19 related models with guidance from this study.

4.5 CONCLUSION

Delineating medical makers and their activities from maker culture helps to describe the infrastructure required to ensure safe and verifiable artifacts. Participants detailed several pragmatic approaches to support medical making, despite risks and challenges inherent to medical practice. Regulatory responsibility being placed on clinicians who make is incongruous with the distribution of labor and skill among medical makers. Clinicians assume nearly all of the responsibility for mitigating risks but have the least facilities for evaluating risk, while engineers and open-source contributors have the most influence over designs and their consequences. Ideally, policies could be adopted that distribute this burden across stakeholders. The engineers interviewed are ready and willing to accept this responsibility as demonstrated by their extensive licensing and documentation practices that go beyond standard maker activities. However, the regulatory infrastructure does not have mechanisms for them to accept responsibility alongside their clinical colleagues. Clinicians alone rarely have the expertise or resources to evaluate these consequences effectively, and rely on these engineers to support them. However, without a change to regulatory policies and institutional practices, this mismatch of responsibility and capabilities will continue to stifle medical making.

There is an opportunity for advances in clinical CAD tools to reshape the distribution of labor among medical making stakeholders. While clinicians will likely retain the ultimate responsibility for what is given to patients, tools can limit what clinicians deploy to only the most safe, tested, and documented designs. By creating CAD tools that enable makers to document their design goals and evaluations while designing, this information can be delivered to clinicians at the point of care. Further, medical-maker, open-source repositories need to do the work of collecting and reviewing designs and deliver them in a consistent fashion to point of care medical makers.

This work captures a moment of medical making that has since passed. Prior to the COVID-19 pandemic, medical making was relatively rare and consisted of many separate communities. While an emphasis on safety and regulatory gaps remain, the ecosystem has undoubtedly evolved as communities of hundreds of thousands of makers stepped into healthcare infrastructure to fight the pandemic. This presented an exceptional opportunity to put our guidelines for a medical maker repository to the test. Following studies of medical making during COVID-19 revealed similar themes, but also demonstrated more variance across medical makers; both what they value and how they uphold those values.

5 A MEDICAL MAKER'S PLAYBOOK: SAFETY-CRITICAL MEDICAL MAKING FOR COVID-19

The prior chapter raises questions about the safety-critical design practices of medical maker communities. The large and unprecedented response of medical makers to the COVID-19 pandemic offered a unique opportunity to examine these practices as communities formed, made medical devices, and delivered them. In response to the COVID-19 pandemic, makers collectively designed, produced, and distributed PPE to assist healthcare and front line workers. In particular, a community of over 2,000 self-described makers from Colorado, Make4COVID, formed in mid-March and had delivered over 80,000 pieces of 3D printed and hand crafted PPE by the end of June 2020 to clinics and hospitals across the state. Makers produced the PPE independently in their homes and organized a unique, distributed supply chain to collect and distribute devices. To do this, Make4COVID enforced a safety-focused approach to making, with the goal of doing more good than harm.

Based on an ethnographic study of Make4COVID, we examine the barriers Make4COVID had to overcome in their efforts to make safely. Make4COVID often struggled to exchange design information between clinical experts, design teams, and makers. This slowed the design of critical PPE that required clinician input, and led makers to produce lower quality prints, which could have increased risks. The community overcame these challenges with the efforts of key facilitators who relayed information between design teams, clinicians, and makers.

We argue that the key limitations of most CAD tools is that they focus solely on the work between a single maker, usually a technical expert, and their machines. Medical making, by contrast, is collaborative and requires makers, clinicians, designers, and engineers to work together to design and make safe products. Clinical CAD tools must adapt to the various forms of expertise that are available in diverse maker communities and amplify their unique contributions. Further, the products of these tools must be portable across varied medical contexts: individual 3D printers, limited resources, smaller teams. This would enable medical makers to produce safe and acceptable designs at scale.

This chapter presents the rich history of a medical maker community who responded to an unprecedented crisis. Based on Make4COVID's practices and the disruptions in medical making activities, we argue that design tools could better facilitate the participation of diverse groups of makers by articulating how design features relate to a designer's intent.

5.1 METHODS

5.1.1 *Engagement with Make4COVID*

We present an ethnographic study of Make4COVID, conducted between March 23rd and September 1st of 2020. I participated as an organizer for Make4COVID. During the acute phase of the U.S. outbreak of COVID-19 (i.e., mid-March through mid-May), I worked between 40 and 60 hours per week inside the community. Hourly commitments slowly dropped off to between 15 and 20 hours a week by late-July and before exiting the community at the end of August. Additionally, I interviewed 26 core organizers summarized in Table 4. During this period, the research team was engaged in multiple studies: surveying other online maker communities [105], interviewing key intermediaries between clinics and maker groups [146], and analysing outcomes of the NIH 3D print exchange. These studies and the broader context of medical making during COVID-19 informed our interpretations of Make4COVID's activities.

I entered the Make4COVID community through my high-school engineering teacher, who was 3D printing face shields for Make4COVID. After an introduction, core organizers invited me to join the community, where I transitioned between different roles: a volunteer (mid-March to April), an organizer (April to May), and a core-organizer (May to September). I advised on community design projects, maker community practices, and quality control measures. By the end of the summer, I was positioned as the leader of

Table 4: Interview participant demographics. Some demographic data was not reported at the request of the interviewee and is marked as not-reported (NR).

Pseudonym	Community Role	Gender	Age	Profession
Isabel	Warehouse Lead	Female	22	Undergraduate Student
Annie	Finance Lead	Female	32	Finance
Shannon	Lead Organizer/Founder	Female	36	Organization Facilitator
Rose	Scientific Advisor	Female	37	Assistant Professor of Biochemistry
Mary	Lead Organizer/Founder	Female	44	Associate Professor Department Director
Claudia	Social Media Lead	Female	46	Community Engagement Coordinator
Janet	Sewing Community Manager	Female	48	Engineering Project Manager
Jennifer	Needs Assessment Lead	Female	53	Self-Employed
Ashley	Regulatory Advisor	Female	NR	Regulatory Compliance Expert
Jessica	Soft Goods Design Team Lead	Female	NR	Industrial Designer
Joshua	Design Team Lead	Male	27	Clinical Design Engineer
Mathew	Lead Organizer/Founder	Male	30	University Lab Manager
Anthony	Thermometer Design Team Lead	Male	30	Military Officer
Ted	Strategy Lead	Male	32	Founder/CEO of Medical Device Startup
Jacob	3D Printing Team Lead	Male	32	Automotive Dealership Parts Counter-person
Robert	Technology Team Lead	Male	41	Associate Professor of Computer Science
Joe	Design Teams Lead	Male	43	Industrial Designer
Michael	Warehouse Lead	Male	44	Air Force Instructor Pilot
Amir	External Engagement Lead	Male	46	Marketing Director
Nick	Scientific Advisor	Male	51	Associate Professor of Biochemistry
Bill	Supply Chain Lead	Male	58	Software Implementation Manager
Ethan	Clinical Interface Lead	Male	62	Mechanical Engineer
James	Floater	Male	76	Retired IT Professional
Kevin	Design Teams Lead	Male	NR	Industrial Design
Liam	Design Team Lead	Male	NR	Mechanical Engineer
Sam	Regulatory Advisor	Male	NR	Medical Device Regulatory Consultant
Gabe	Community Stewardship Lead	NR	NR	Web Developer

a new team dedicated to writing the "Make4COVID Playbook". The *playbook* was a short series of essays describing Make4COVID's decision making processes and presenting anecdotes in a format that could support similar maker communities. The playbook was intended to provide other maker communities or crisis response organizations insight into Make4COVID's successes and failures.

5.1.2 *Data Collection and Analysis*

During this study period, I took field notes, wrote memos about salient experiences, collected Slack and email conversations (as well as channel archives), reviewed public community forums (hosted over Zoom and stored on YouTube), and interviewed 26 core organizers. Due to the large amount of data, only salient emails and threads were included in initial rounds of thematic analysis. Other content was archived and revisited as different themes were explored. We interviewed every core organizer who attended a stand-up meeting between June 1st and September 1st and every design team lead, even those who were no longer regularly attending stand-ups.

Each week, the research team met to discuss the week's events, following Lincoln and Guba's structure for debriefing sessions [51, 155]. During these meetings, we conducted a thematic analysis [179] where the most salient notes and artifacts from the week were inductively coded. Following guidance from Fine [70], we examined Make4COVID as an ongoing culture centered around shared virtual spaces, interpersonal relationships, and a shared history. To contextualize that shared history, researchers brought contemporary media coverage of global COVID-19 Maker efforts, as well as notes from contemporaneous studies [105, 146]. This outside data was used to triangulate [51] Make4COVID's position in the wider context of the COVID-19 pandemic and compare its actions to those of contemporary groups. Emergent themes were discussed and noted for future review. Each week, these themes and codes were iteratively revised. Finally, to ensure the validity of our emergent themes, we conducted a form of member-checking [51] by presenting drafts of the playbook to various organizers. This gave members an opportunity to correct our findings.

5.2 BACKGROUND ON MAKE4COVID

"Make4COVID is a coalition of volunteers designing, manufacturing, and distributing essential equipment for Colorado's health care workers and first responders. We come from all walks of life. We are makers, designers, artists and engineers, hobbyists and professionals, from across the state of Colorado and beyond, united in common purpose. " (Mission Statement)

Make4COVID is a coalition of self-identified makers that developed from a group of colleagues from the University of Colorado (CU) Inworks program [120]. In collaboration with the CU Anschutz Medical Campus, the team began organizing a rapidly growing community of volunteers to 3D print and distribute PPE to frontline workers across Colorado.

The state saw a slow and steady rise in positive COVID-19 cases and hospitalizations from March to August 2020. The state was put into lockdown from mid-march to late may, where all residents who were not essential workers were asked to stay home. The large number of people staying home became a valuable source of volunteers. Prior to the pandemic, the state had made significant efforts to attract manufacturers and tech companies to bolster the local economy. This growing industry proved fruitful during the pandemic [93], particularly in the Denver, Front-Range area. The tech infrastructure created by these industries significantly contributed to Make4COVID's resource pool (e.g., donated materials, skilled volunteers, funding).

By the end of March, the team had grown to include 30 core organizers, 271 organizing members, and 2,218 volunteers. By mid-June, the community had collectively delivered 81,532 pieces of PPE , in the form of face shields and sewn cloth masks. Since closure of the study, that number has risen to more than 120,000 piece of quality-controlled PPE. This PPE was distributed across Colorado, with priority given to rural hospitals and clinics that could not access traditional PPE. The majority of makers were based in the densely populated Front-Range, but community members lived in every county, with some even living out of state. Warehouses were centered in Denver and Colorado Springs, and remote regions were supplied by volunteer pilots from Colorado's Civilian Air Patrol [233].

The 81,532 pieces of PPE delivered represents only a fraction of what Make4COVID produced during the spring and summer of 2020. Particularly, early on, a large amount of what makers produced was thrown out due to quality control measures at warehouses. Further, one team of volunteers was dedicated to checking in with clinics that received the PPE, and identifying any unexpected outcomes or risks. Despite this long term investigation, no clinics reported any failures in the delivered PPE. Further, many clinics reported back usage of face shields and masks as late as August 2020, five months after receiving their first shipments. It is possible that broken or flawed designs were delivered, but the lack of reporting is strong evidence that if this occurred, it was rare. Unfortunately, there are no studies, either by Make4COVID or health officials in the region, that provide strong evidence of the efficacy of this PPE at preventing COVID-19 among healthcare professionals. This means that while the good Make4COVID did is unknowable, their devices did minimal or no harm. That is, Make4COVID made safely.

5.2.1 *Organization and Communication Platforms*

Make4COVID quickly developed a top-down organizational hierarchy. Core organizers directed community policy, goals, and efforts by forming specialized teams. A wider group of organizers in these teams managed specialized activities: design, supply-chain, media relations, legal, finances. These teams relied on a wider pool of 3D printing, sewing, and delivery volunteers. They contributed as much or as little as they chose, by repeatedly completing specific tasks independent of other volunteers. For instance, makers would 3D print face shield headbands at home. Sewists would sew cloth face masks and head straps. Drivers would go to pick up locations, collect makers' and sewists' products, then deliver them to a regional warehouse. Volunteers received guidance from the community on how to conduct their work safely. Safe practices were defined by teams of experts (e.g., clinicians, biomedical engineers, health policy makers, virologists) based on forthcoming scientific consensus about the spread of COVID-19. Beyond these guidelines, they operated with significant autonomy.

There were three key platforms that community members used to communicate: Slack, Mighty Networks, and the sewist phone bank. The Slack work space was originally created to organize the community, but as the community grew, it became cumbersome to manage differing channels. Within weeks, the Mighty Network was formed and the majority of members were pushed to that platform to gather instructions about what needed to be made and how to make it. As the community started efforts to sew masks and head straps for the face shields, a separate system formed among sewists to communicate through a phone bank and email lists. Over time, the separation in roles of members of each platform became more apparent.

Participation in Make4COVID was open to the public through a Mighty Networks [29] forum linked to their web page. This was where most volunteers collected information about available tasks, built relationships with other members, and provided feedback to organizers. Early on, volunteers tended to enter the community based on word-of-mouth, but as Make4COVID's efforts were more widely publicized on social media and on local news broadcasts, more volunteers joined directly through the community website.

Not all volunteers readily adopted the Mighty Network because it did not fit sub-communities' mechanisms for communicating. The sewist team lead, Janet, developed a "*hub and spoke*" phone banking network to communicate with sewists. Janet managed a team of "*hub captains*" in different local communities around the state. Those hub captains then held conference calls with sewists in their area, delivering daily updates, receiving feedback, and maintaining social connections that acted like a form of socially-distant sewing circles. The content of these phone calls was supplemented by email lists which would disseminate Standard Operating Procedures (SOP), sewing patterns, and other key pieces of information. In her interview Janet reflected that this structure emerged to fit the needs of the sewists, a group that had more women and people over 65 years of age than the overall Make4COVID community,

> *"Women are social beings, we also had a lot of folks who were older. They wanted some way to give back. We gave them an opportunity by meeting them where they are. That required some one-on-one personal connection. They wanted to connect with us. I think the communities are really quite different. We have a good number of sewists who didn't go to the Mighty Network. We had to bend to a number of different communication protocols. " (Janet)*

Organizing teams were separate from the wider volunteer community, and collaborated using an invite-only Slack work space [263]. It consisted of 75 public channels, as well as an unknown number of private channels, where members discussed channel topics asynchronously. Most channels represented individual teams, however some teams broke up their discussions across multiple channels. Teams individually managed their activities. Large teams held daily video conferencing stand-up calls. Smaller teams met as needed, and relied on personal relationships to structure their work. As teams completed their work and dissolved, most members would stop contributing to the project, though some would seek out new roles through their team leads.

Finally, the small set of core organizers was primarily made up of community founders, team leads, and experts identified by other core organizers. In addition to a private Slack channel, the core organizers met daily for morning stand-up meetings between March and May, dropping down to three meetings per week during June, and finally ending with two meetings per week in July and August. During these meetings, core organizers discussed Make4COVID's mission and facilitated cross-team collaborations.

5.2.2 *Make4COVID's PPE*

Make4COVID conducted a variety of design efforts to produce different types of PPE. Early on, the community focused on face shields, sewn masks, and ear-savers (Figure 4) at the request of hospitals. Instead of starting from scratch, the face shield design team modified the pre-existing Prusa face shield [206], which had gained recent popularity among clinicians. With the support of clinicians at the CU Anschutz medical campus, the design was adjusted so that it could be more easily printed on a variety of consumer 3D printers and quickly released to the maker network. These types of PPE were the only ones produced and delivered by Make4COVID. Notably, all PPE designs that were delivered by Make4COVID were tested using the same standards as equivalent FDA-approved medical devices. For example, face shields went through the same testing as traditional disposable face-shields, while straps, which have no FDA-approved equivalent, were put through stress-tests created by design teams of sewists, engineers, and clinicians.

Aside from Make4COVID's main production efforts, teams set to work from early March through the late summer to design more advanced forms of PPE (Figure 5). These types of PPE required a more extensive review process to align with regulatory standards. For example, the N95 respirator designs were put through the same testing process as traditional, FDA-approved N95 respirators. The design teams were unable to come up with workable solutions that could use materials and manufacturing methods available to the community and still uphold these rigorous safety standards.

By mid-summer, community design projects were instigated, in response to calls from makers to have more of a role in the design teams. On April 19th, Joshua, Gabe, Ted, and Joe reached out to me to instigate a community design effort. Joe and the other design team leads selected relatively low-risk projects that they would open up to the makers. On April 29th, they put out a call for designs for a contactless thermometer and reusable, sewn medical booties (Figure 6). Once the designs were selected, the teams were formed from applicants. These projects were considered low risk because they were less likely to cause the user or patient injury if they failed (e.g, a PAPR hood providing air to the wearer vs a booty over the shoe), not the wearers primary protection against the virus (e.g., an N95 mask vs. a face shield), and not used in high risk clinical activities (e.g., intubation vs. checking a clinic visitor's temperature). The new production efforts became closed off to the community, just like the design efforts.

5.2.3 *Making and Delivering PPE*

Once a design was established, a call was put out to makers through the Mighty Network and sewist phone bank to make the design for distribution to clinical facilities across Colorado. Information about the designs was available through the Mighty Network's "Start Making" page and a sewist email list. The designs included the 3D models for 3D printing, sewing instructions, recommended print settings, the SOP, and PPE drop off instructions. From here, volunteer drivers gathered maker's drops, sanitizing the stations after collection, and brought them to one of the two warehouses. Warehouse volunteers sorted through deliveries, reviewed and tested the PPE, took inventory of the results and the individual makers' records, sanitized the products, and packaged them for delivery. A team of Needs Assessment organizers maintained a list of requests across the region, prioritized by clinical relevance (e.g, hospitals, clinics, dental offices), access to traditional PPE (e.g., urban hospitals, rural clinics), and past deliveries. Once an order was filled, drivers would collect shipments to deliver to clinics. In the case of distant, rural clinics, volunteer pilots with the Civilian Air Patrol delivered the shipments.

Quality assurance processes evolved substantially in the early months. Initially, Make4COVID had no official quality control policy. Diligent warehouse volunteers chose to check prints without instruction and raised concerns about quality with organizers. Early estimates of the portion of failed devices were as high as 60%. This lead the teams to develop a quality control protocol, which makers were instructed to follow before sending in their PPE. This checked common errors like sizing and print quality. After an adjustment period, makers began to send in much higher rates of quality designs. Warehouse staff continued to check each piece of PPE while sanitizing and packaging them.

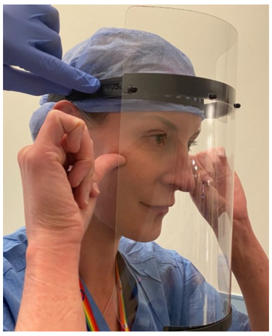

(a) Face Shield

(b) Sewn Cloth Mask

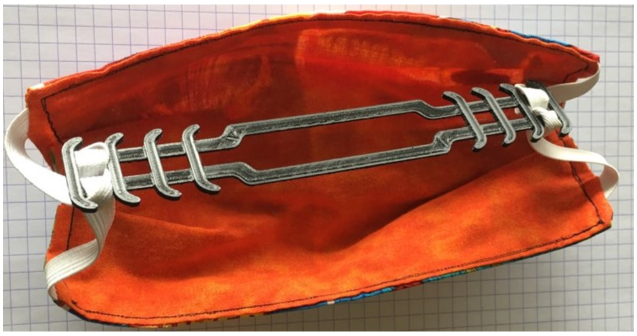

(c) Hold

Figure 4: Sample images of Make4COVID's class 1 maker-made PPE

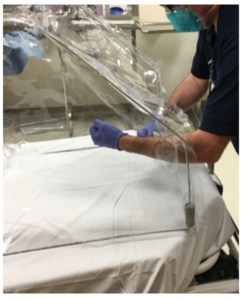

(a) Intubation Shield

(b) N95 Respirator

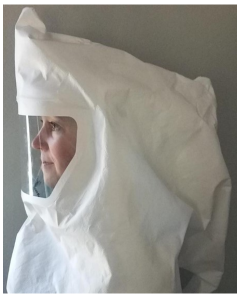

(c) PAPR Hood

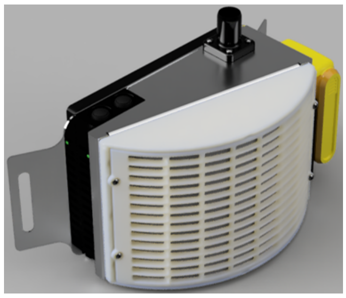

(d) PAPR Pump

Figure 5: Sample images of Make4COVID's Higher Risk Design Efforts

(a) Protective Shoe Covers, aka Booties

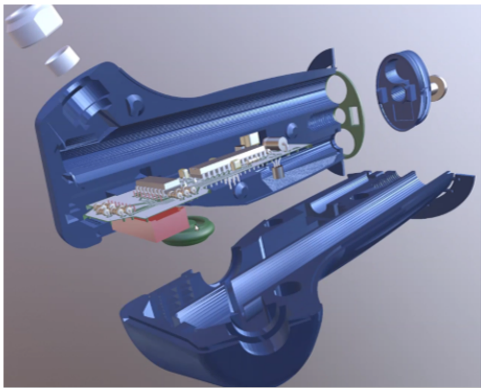

(b) Expanded View of Thermometer Model

Figure 6: Sample images from Make4COVID's community design projects.

5.3 RESULTS

This study revealed three salient themes related to Make4COVID's safety-oriented making practices and how they affected the broader participation of diverse community members.

Open and Safe Making: Make4COVID's separate design and production activities comprised multiple efforts to include makers in safe practices.

Formats of Medical Making: The affordances of design formats (e.g., 3d models, documents, sewing patters) affects who could easily collaborate and utilize Make4COVID's designs.

Facilitating Collaboration across Domains: Facilitators were critical in ensuring broader participation of clinicians and makers.

5.3.1 *Open and Safe Making*

Making safely is a principle concern of medical making, and prior work [105] has shown that medical making communities often struggle to balance the goal of including a diverse set of makers with ensuring the quality and safety of their products. By segmenting Make4COVID's efforts into two types, design and production, Make4COVID offers two visions of what is needed to make safely and who can participate in that process.

5.3.1.1 *Closed Designing*

Make4COVID produced a variety of design artifacts including: SOPs, 3D models, assembly guides, and regulatory briefs. Once design efforts were completed, these artifacts were used in production efforts, where the community produced and distributed as much of the designed product as the community needed. While production of face shields, sewn cloth masks, and ear savers was inclusive of a large and diverse set of makers, design efforts were largely managed by small teams with professional design and engineering expertise. Within the community, this raised questions about what open and inclusive entailed; is a community open if only production, but not design, includes all makers?

One view among members of the design teams was that, while more open design efforts could have value, they were not compatible with safe making practices and would disrupt and delay activities. The design teams lead, Joe, firmly decided that the design teams would not release their work before it was completed. He was concerned that allowing public access to incomplete or unverified designs would pose a risk, as uninformed users could reproduce it without understanding the failures possible in the designs. From the design team leads' perspectives, safe design was the priority, while community engagement was an expendable benefit. One lead, who shared Joe's sentiment, put it bluntly,

> *"When we selected the proposal, the community outreach objective was met. Our primary objective was to create a functional product. " (Anthony)*

Anthony is noting the implicit ranking of values shared by the design teams. Openness and inclusion of makers remains a value, but is secondary to safety (i.e., "*a functional product*"). He implies that community outreach was only needed to source ideas and designers.

However, outside of design teams, other organizers raised debates over how Make4COVID would open source their work. While Make4COVID's mission statement included intentions to open source all of their designs, the primary debate was about how open the design process should be in practice. Advocates within the community pointed to groups like e-NABLE as hallmarks of open source medical making. One organizer encouraged Shannon to pursue a stronger relationship with a Colorado e-NABLE chapter:

> *"I've printed parts for [e-NABLE] and other support in the past. They are legit open source too... if you want to see a case study in open source failure, compare "openbionics" with "e-NABLE". Openbionics was just open source in words and promotion, but didn't deliver on it... e-NABLE actually did do everything with open source methods and licensing, and their hands are in use around the world. " (Slack General Thread)*

During one core organizer stand up, which focused on this debate, multiple design team leads contested this organizer's claim that e-NABLE produced safe designs, citing the lack of studies of the long term effects of their prosthetic-like devices, and the minimal efforts by the e-NABLE community to provide safety guidelines to makers. Consensus could not be reached on whether larger open-source teams could produce safely or how this could be managed on Make4COVID's short, six-month design period. Ultimately, this organizer left Make4COVID in mid-April. While safe and open design may be possible, Make4COVID could not find a satisfactory solution that also fit into the time and resource constraints of an ongoing crisis.

Concerns about makers on the Mighty Network being unable to participate in design caused strife between makers and organizers. Many of the makers who joined Make4COVID demonstrated professional and technical expertise, and expected to bring that expertise to bear by contributing to the design projects. When design roles were not readily available, organizers would push these volunteers towards PPE production. One design team lead, Joshua, reflected on this challenge in his interview:

> *"We were kind of losing members because they come in and say, 'I have 30 years experience doing this. I can do that at a very high level'. And then we say, 'Great. Here's a 3D printing file, go print'. We could tell that was a little deflating for a lot of people. " (Joshua)*

These ongoing conflicts between open and safe design practices demonstrate a broader challenge for participation in medical making. There was a common desire to include more makers in design activities, however, concerns about their ability to create safe designs necessitated an ongoing review of any designs that came through. Without this review, the clinicians who received these devices would be responsible for evaluating their safety. This conflicted with Make4COVID's stated goal of reducing the "*strain*" on hospitals responding to the the COVID-19 crisis.

> *"We're taking the strain off of already busy hospitals by giving them a single point of contact for supplies coming from Colorado makers. " (Mighty Network FAQ)*

Alternative models were proposed (e.g. e-NABLE), but technical and clinical experts disagreed with the assessment that those communities followed sufficiently safe practices.

5.3.1.2 *Open Manufacturing*

While participation in design efforts was largely restricted to technical experts, Make4COVID's production efforts were open to anyone through the Mighty Network and sewist email list. Once designs were created, production could be broken down into tasks which were accessible to a variety of volunteers. Volunteers with access to 3D printers could print face shield headbands and ear savers. Experienced sewists could make masks, while more novice sewists could focus on easier head straps. Volunteers without crafting and making expertise and resources could support delivery to and from warehouses. An internal community survey conducted in June 2020 showed that 1609 of the then 2159 member community had contributed by making PPE (e.g., 3D printing, sewing). The scale of that membership demonstrates the openness of community's production efforts.

Multiple organizing efforts were aimed at keeping the production efforts inclusive of a diverse maker community. Within weeks of starting face shield production, the supply chain organizers had developed a system of delivering PLA filament and sewing supplies to makers so that the ability to pay for or acquire materials did not dissuade makers. Similarly, the supply chain team managed a large network of drivers and drop off points across Colorado so that makers with limited transportation could gather materials and deliver their PPE. Subgroups within the community created support systems for makers with accessibility concerns. For instance, the sewist team used phone calls and emails rather than Mighty Network, because the Mighty Network was not as accessible to the older sub-community. Each of these efforts demonstrates the organization's commitment to broadening participation of makers in their production efforts. These efforts meet utilitarian (e.g., increased supply) and cultural (e.g., inclusion in a community act of care) prerogatives.

5.3.2 *The Formats of Medical Making*

The formats (e.g., 3D models, sewing patterns, external documents) Make4COVID used to present information revealed what types of information can be expressed in the design tools they adopted. To examine this aspect of collaborative making, we review the types of information that were included in Make4COVID's PPE designs, the formats that presented this information, and how they affected makers' outcomes.

Consider face shield designs; beyond the 3D model, the design includes safety protocols for maintaining a sanitary work space, print settings for ensuring the final product's material properties, and instructions for evaluating the print's quality. Each of these components was presented to makers in different formats. The shape of face shields was given through STL files that exclusively showed its form and size but provided no information about use or manufacturing. SOP documents provided the needed manufacturing and safety instructions. Print settings were provided in a list of details in the SOP and on the website, but these often needed to be adapted to each unique printer. Some details about material properties (e.g., flexibility, fit) were never documented outside of design team's Slack conversations. Additional details about use and disinfection were included in an Information for Use (IFU) document sent to clinician recipients. These are different than SOP's, which were given to makers. SOP's provide manufacturing details, while IFUpackets serve as a safety-focused user manual. The presence of so many documents describing the face shield designs demonstrates that 3D models alone cannot express critical information about the design.

Unfortunately, the separation between the models and the documents that provide critical information about their manufacturing and use, led some makers to 3D print PPE unaware of these instructions. To combat this, makers had to check a box acknowledging they had read the protocols posted on the Mighty Network before getting assigned a drop off location. Unfortunately, this mechanism could not guarantee makers followed protocols, let alone understood them, and relied heavily on community trust. Further, this mechanism only enforced gathering instructions once, rather than each time they were updated.

Beyond critical, functional information about designs, these external documents were a critical communication tool to establish trust with makers. Initially, the SOPs provided only instructions about the tedious process of printing safely (i.e., reducing viral contamination of the printed object), with little explanation of why these steps were necessary. In response to these inexplicably lengthy protocols, one maker protested by sending in a bag of dirty paper towels, labeled "SOP-Proof" (Figure 7a). They had used these paper towels to clean their print space and wanted to definitively demonstrate that they had followed the lengthy protocols. In response, the SOP FAQ was updated to include an explanation of how COVID-19 could be transferred from the maker to the print and then to the warehouse volunteers or clinicians that handled it. However, adding these details to the document was not sufficient, because makers could easily miss updates to these documents.

Just as 3D models could not convey sufficient information about the design, makers had to adapt print settings to fit their personal work space, and this led to critical quality concerns. A low quality print may not fit with other face shield components (e.g., the clear plastic shield) or break at key points (e.g., the stub that connects the shield and head band, the flexure of the head band that conforms around the wearer's head). Small differences in size, layer adhesion, or print density could lead to critical failures (Figure 7). Prior to instituting a quality control rubric, Isabel, a warehouse organizer, estimated that early batches of face shields included 40-60% failed prints that included these common errors. During her interview, Rose reflected that her team had little practical experience with 3D printing, and had expected prints to be ready for distribution without much quality control:

> *"We hadn't really gone that far down the rabbit hole. We didn't really know how advanced the community was with printing stuff and we hadn't really done any of our due diligence there. " (Rose)*

This demonstrates a disconnect between the consistency design teams expected of makers, and the reality of making in so many unique contexts. The expert design teams had a mental model of what a 3D printer could consistently produce, but had not considered how small differences between printers could affect quality.

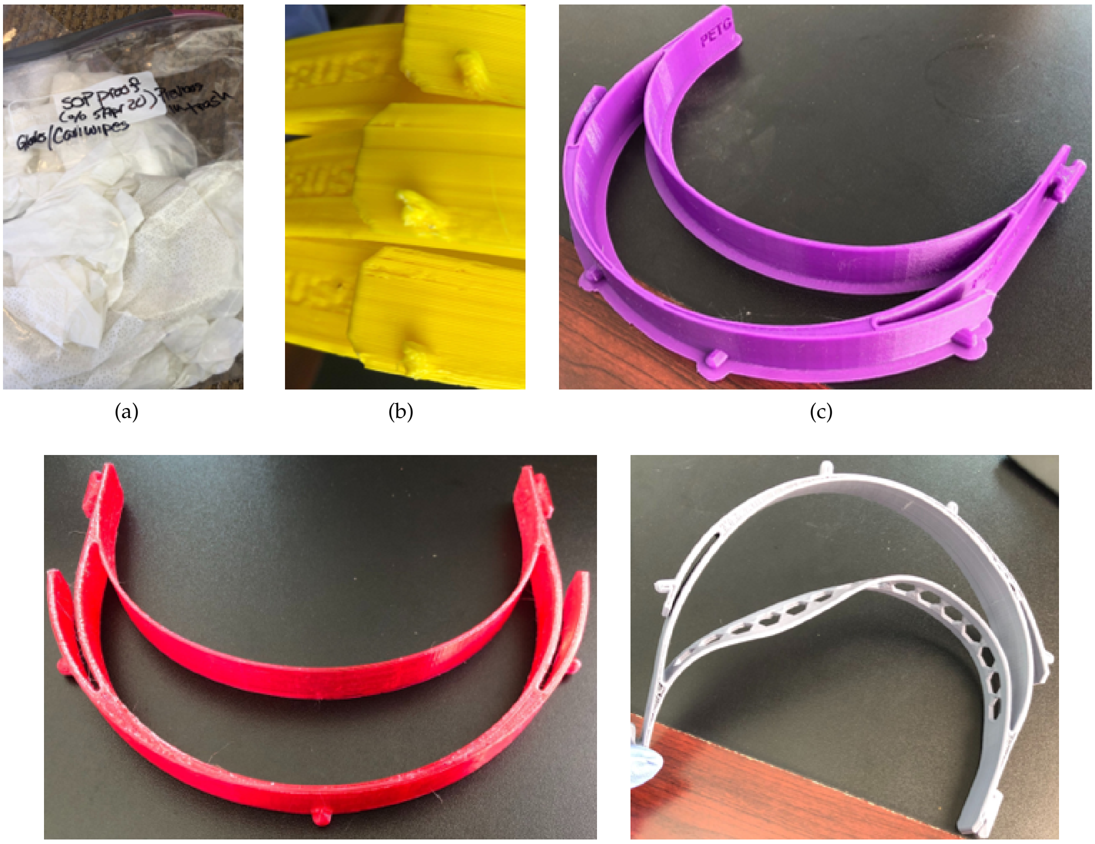

(a) (b) (c)

(d) (e)

Figure 7: Images taken from quality control Slack channel, showing face shields that failed warehouse quality control because of: (a) inclusion of hazardous dirty tissues, (b) delaminated and malformed mounting pegs, (c) use of a brim (to connect to the 3D printer bed) violating the SOP, (d) a print of a similar but invalid model, (e) dramatically bent and unusable product.

To ensure quality, makers needed information about how to evaluate their designs. In response, Ted and I introduced new quality control protocols that makers could follow at home. This quality control rubric helped makers evaluate their prints. It is noteworthy that the team could not find a way to embed this information into the design models or printer settings, instead providing the quality control rubric as a mechanism to support makers learning to adapt their work flows to the community's standards. This method was not useful in isolation; many inexperienced makers still relied on 3D printing experts like Jacob to tailor their print methods.

In contrast to the face shield designs, Make4COVID's sewing patterns were streamlined and condensed to one document. The sewing patterns followed the conventions of common sewing patterns from crafting communities, likely because key designers (e.g., Janet, Jessica) were recruited to Make4COVID from these communities. Sewing instructions were embedded in the pattern's form, such as fold lines and captions for seam allowance (Figure 8a). Sizing was conveyed through scaled printouts which sewists could use to measure their work (Figure 8b). Instructions were embedded with explanations of key design decisions and exemplary photographs (Figure 8c). Finally, whimsical images and words of encouragement urged on volunteers and complimented their efforts (Figure 8d). 3D printing workflows only rely on 3D models and print settings, which let makers accidentally or intentionally ignore external protocols. By contrast, sewing is a manual process that requires the sewist to interpret the design at each step. As such, the sewing patterns included all of the relevant protocols alongside the form of the design.

The formats Make4COVID used to convey its designs to the community are commonplace in the 3D printing and sewing domains. They also contributed significantly to the quality of the PPE. When instruc-

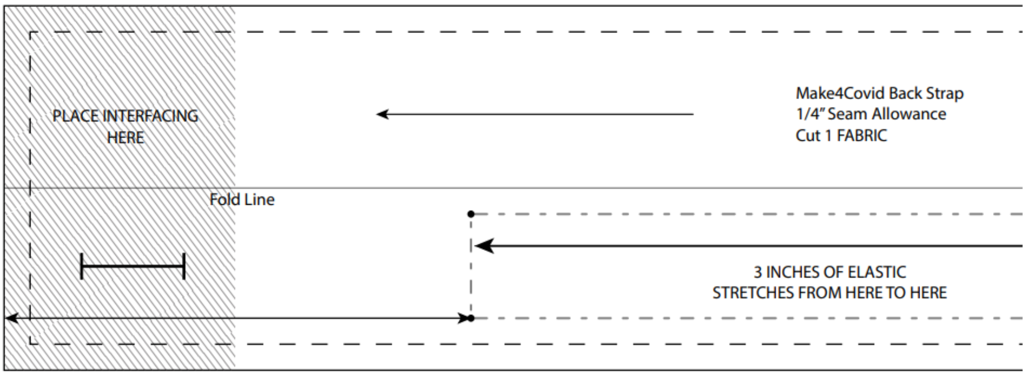

(a) Instructions, such as seam allowance and fold lines, embedded in the pattern design directly

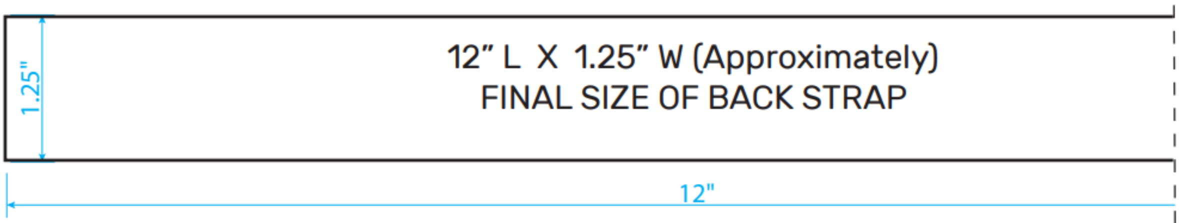

(b) A scaled pattern for measuring fabric components

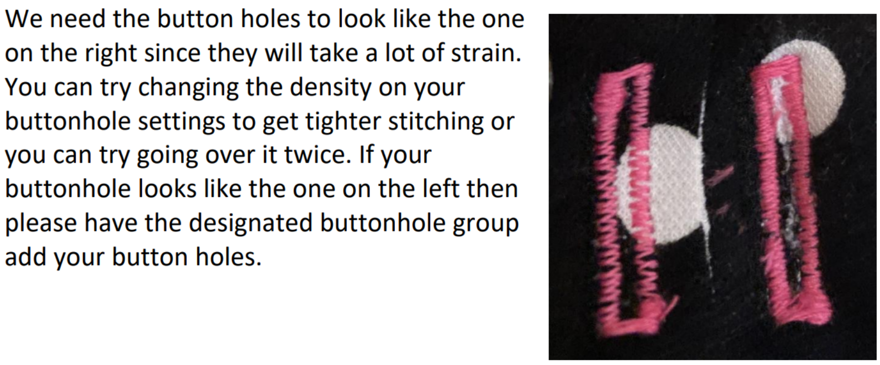

(c) Details about the density of stitches and its purpose beside exemplary images.

(d) The instructions end with this whimsical teddy bear and words of encouragement such as "masterpiece".

Figure 8: Examples of instructions and design details in the back-strap sewing pattern

tions and design decisions were isolated from the 3D model, it confused makers and led to lower quality products. Despite recognition of these limitations, designers of the face shield could not fit this information into one format. By contrast, the pattern design sewists adopted from the wider craft community enabled designers to inject key information throughout the pattern document, and may have resulted in a higher incidence of quality production. It is worth noting that sewist communities, and crafting communities at large, adopt generalized formats (e.g., Microsoft Word documents, drawings) to share their designs, while 3D printers use specially tailored CAD tools. This gives crafts communities the flexibility to tune their design formats to what helps crafters most, rather than to software constraints. Design tools for 3D modeling may necessarily be constrained to inter-operate with a variety of devices, but, when designing these tools, the practices of craft communities are a rich source of templates and best practices.

5.3.3 *Facilitating Collaboration Across Domains*

Make4COVID consisted of multiple sub-communities (e.g., organizers, design teams, clinicians, sewists, makers), and facilitators were critical to connecting these groups and enabling collaboration. We discuss two key types of facilitators. First, "community stewards", who bridged the gap between organizers and the makers, both on the Mighty Network and in the sewist community. Second, clinical facilitators, who organized interactions between Make4COVID's clinician constituency, the design teams, and core organizers. Without these facilitators, Make4COVID would have limited its participation to technically-oriented teams of industry professionals, rather than encouraging the participation of amateur makers.

5.3.3.1 *Connecting Organizers and Makers*

"Community stewards" were critical members of the community that resolved makers' confusion about the designs and protocols. Recall the confusion among makers fostered by the separation of key manufacturing instructions from the face shield 3D models. Recommended print settings had been provided on the Start Making page of the Mighty Network, and many makers in "3D Printer Help Zoom Calls" reported using them. However, these settings often needed to be adapted to differences in printers or filament brands. "Community steward" Jacob led an effort to provide additional settings and instructions for common printers by sourcing solutions from makers who attended the Zoom calls. While useful, this proved an intractable solution because there were too many different printer setups across the community, each of which needed its own settings and instructions. Meanwhile, makers could not reason about the quality of their prints. These challenges likely dissuaded many makers from participating and created difficult barriers that others worked diligently to overcome.

The "community stewards" were critical facilitators that maintained Make4COVID's social cohesion. For instance, Jacob's '3D Printing Help' Zoom calls enabled a core group of makers to socialize and discuss printing challenges. Jacob also facilitated the formation of a pool of 3D printing experts, the "3D Print Crash Testers", who were a critical resource for design teams who needed prototypes made. In this group's calls, highly-engaged makers in the community would debug each other's printers and develop guidelines to help other makers adapt the community's designs to their work spaces. Without Jacob running these, at times, 12 hour long calls, Make4COVID could not have sourced feedback on face shield printing instructions, which proved critical for creating quality control measures. It also created a mechanism for design teams to identify domain experts among the Mighty Network and bring them into the design teams. Without facilitators like Jacob, and the social network he fostered among makers, large amounts of community expertise would have been ignored.

5.3.3.2 *Relaying Clinical Expertise*

Safety was one of Make4COVID's core values, and they relied on clinical and regulatory experts to ensure their making practices would do no harm to the end users. While medical experts were among Make4COVID's ranks, enacting their feedback and design requirements was a constant challenge in all design activities (e.g., low and high risk).

Once the scope of Make4COVID's PPE delivery exceeded the CU Denver network, Jennifer, Joshua, and Ethan formed a clinical-needs assessment team to facilitate an ongoing exchange between organizers, design teams, and the clinicians who received the PPE. Joshua facilitated the collaboration of engineering teams and clinical teams out of CU Denver during Make4COVID's earliest design phases. As Make4COVID expanded, Jennifer managed relationships with recipient facilities, and regularly checked in on how the PPE was being used and if any challenges arose. Ethan worked with a subset of recipients to gather expert reviews and feedback about Make4COVID's design projects, and then relayed it back to the relevant teams.

A physical prototype, particularly of high-risk devices, was critical to gathering feedback from clinicians, but this created a significant burden on design teams who wanted to iterate on lower-fidelity prototypes (e.g., sketches, 3D Models). Both Joshua and Ethan noted that clinicians would rarely provide feedback without a functional prototype to test out. For instance, Ethan remarked on the requirements Emergency Room doctors had for reviewing an intubation shield design,

> *"This doc has been so helpful in explaining the use cases. Eloquent on the subject. The doc felt that it was essential in at least one or two cases to be present in ER and see the device. [They were] very willing to do this kind of thing. " (Ethan)*

Ethan explained that the clinical feedback tended to be very early in the process, helping explain use cases or provide requirements, or late, when there was a prototype to test. However, clinicians all but refused to review intermediary prototypes, like design sketches or CAD models, because understanding them took more time than they could afford. This often frustrated designers, who viewed building high-fidelity prototypes as a significant and unwarranted burden when low-fidelity prototypes were available. Without clinical facilitators, Make4COVID's design teams could not access the expertise needed to ensure their designs would be safe and effective in clinical environments. However, the tools designers used to iteratively prototype and the level of fidelity clinicians needed to reason about design prototypes were often mismatched. Facilitators stepped in to address these challenges and collect the needed information.

These findings call into question medical maker communities' mechanisms for fostering broader participation and safety. On the one hand, the design practices that draw in many makers are exceedingly difficult to make open. Design teams working swiftly to meet urgent clinical needs cannot provide an open set of makers the resources for proper safety review. Alternatively, technologies like 3D printing, supplemented with sufficient organizational labors, can enable a wide group of people to produce medical devices at scales usually reserved for factories and mass manufacturing methods. The smooth transition from design to production efforts is facilitated by common digital formats (e.g., 3D models, documents, sewing patterns), but each of these afford different critical aspects of the design. Separation of printing and safety information led makers to critical mistakes that affect the quality of their products. Streamlined formats (e.g., sewing patterns) caused less trouble and raised fewer quality concerns. The discrepancies left by segmenting designs across formats were managed by key facilitators between designers, clinicians, and makers. Their efforts are a testament to the volunteers' commitments with Make4COVID, but also raise concerns about how much effort is necessary when insufficient design tools inhibit the participation of diverse community members.

5.4 DISCUSSION

This case study of Make4COVID highlights how medical makers can balance medicine's value of safety, with maker culture's value of openness. In particular, we discuss two groups that can participate in medical making. The first are *orthogonal experts*: people whose primary expertise is separable from and independent of the material expertise associated with making (e.g., craft, design, engineering). If effectively supported, orthogonal expertise enables communities to solve problems in new domains. Orthogonal experts contributions derive not from the designs or products they produce, but from the knowledge that makes them usable. For instance, clinical, biomedical, and regulatory expertise are all types of orthogonal expertise Make4COVID accessed to embed the values of safety and care into their PPE. The second group are the "expert-amateurs" [143] usually associated with making. By isolating this broader group from orthogonal experts, we can discuss where broadening participation to one group inhibits the participation of another. This, in turn, gives us insights into how participation affects the outcomes of medical making, particularly in terms of safety.

5.4.1 *Participation in Prototyping*

We first consider how participation in prototyping and design efforts are broadened; increasing participation of expert-amateurs can promote novelty and empower these makers [245], however doing so may dis-empower orthogonal experts (e.g., clinicians) and undermine medical making values of safety[101], care [250], and end-user empowerment[117]. Making safely, rather than for the "pleasures of production" [245], extends from an awareness of the medical context, where the expected outcome of professional care is to protect the interests of the recipients. In the context of medical making, clinicians embed their orthogonal expertise to mitigate risks, which then empowers them to establish a network of care between the maker and medical communities. As noted in Chapter 3 and Chapter 4, under normal circumstances, clinicians rarely have the time to manufacture devices directly, and, in times of crisis, their ability to participate in any aspect of making is further constrained by emerging priorities for institutional needs [146]. In this case, participation of orthogonal experts does not involve the design or production work of making. They participate by reviewing, critiquing, and guiding design efforts.

Broadening participation to an open network of makers (i.e., "expert-amateurs" [143]), while valuing safety, puts a greater burden on orthogonal experts' reviewing labors. As some Make4COVID members argued, broadening participation in design efforts is possible through open sourcing methods. However, evaluating which designs meet community standards for safety requires the attention of orthogonal experts (e.g., clinicians, biomedical researchers, regulators). Despite careful consideration of this trade-off, Make4COVID could not find a way for makers to openly participate in design without a bottle-necked review by the few clinicians available to the community. While the review burden could be mitigated by facilitators, the added work did not merit the imposition on clinicians. While open participation of makers empowers those makers, that value must be weighed against the burdens it imposes on others. If open participation in the design process remains a priority, the systems that foster design efforts must help amplify the contributions of the few orthogonal experts to review the contributions of the wider network of makers.

5.4.2 *Portability of Production*

One way to amplify orthogonal expertise while broadening participation of expert-amateur makers is to direct those makers to production, rather than design efforts. In this case, we define production as the physical labors of making an established design: printing face shields, sewing masks, cleaning work spaces, assembling products, and sending them to recipients. In this regard, Make4COVID was an extraordinary success, producing more than 80,000 pieces of quality controlled PPE in a matter of months during an unprecedented crisis without any reports of adverse events.

Production of safe designs at scale requires that those designs be *portable* [213] across each makers' work space; that is, *the key aspects of a design that derive from orthogonal expertise must be reproducible by a wide variety of makers*. Make4COVID makers' output revealed the variance a single design can produce across thousands of products made by thousands of makers. The details that affect production by a variety of makers are rarely encoded in the design formats makers can adopt (e.g., 3D models, design briefs). To adapt designs to a new printer or workshop, makers require the embodied knowledge of experts like Jacob, who are familiar with the obscure ways that print settings can affect key properties of the model. No matter how meticulous, how well tested, these designs are limited in their portability to each makers' workflows, which in turn limits who can participate in their production.

In most cases, quality discrepancies across makers derived from small differences between makers' workflows (e.g., Slicers, 3D printers, filaments); these seemingly trivial details are not portable between makers. In other maker communities, where each maker is both designer and producer, they can individually adapt their design to the peculiarities of their work space. When the designer and producer are the same person, portability is insignificant. However, for maker communities to produce consistent designs at scale, singular designs must be portable across each individual work space. Challenges with design reuse and error prevention have been explored before [115, 190, 212, 213, 272], focusing primarily on modifications to existing designs. However, the critical, but often invisible, work of converting a design into a product using a specific set of tools is rarely considered [99].

5.4.3 *Engaging Orthogonal Experts*

In an effort to broaden participation in design and making, prior work often frames makers as existing on a spectrum from novices to experts. This framing influences the technologies maker communities adopt, and by extension creates barriers to orthogonal experts. Technologies aimed at novices simplify the design and modeling process, but limit what makers can easily articulate. Expert oriented tools expand what can be represented, but rely on affordances ingrained in engineering and technical practice (e.g., sketching systems designed like draft boards, visual programming interfaces for managing parameters). Both approaches to design tools require orthogonal experts to adapt their expertise to these tools rather than enabling the tools to adapt to new domains.

When we consider clinicians as orthogonal experts we can see how the affordances of these design tools impairs their ability to embed a value of safety in designs. For Make4COVID, the process of gathering feedback from clinicians was often tedious and full of miscommunications. Clinicians could readily provide design requirements and evaluate a physical and working prototype, but to engineers and designers, minimum viable prototypes include sketches, 3D models, and non-functional or incomplete physical models, which clinicians could not readily interpret. Facilitators between clinicians and design teams worked to ensure that the clinicians' orthogonal expertise was upheld, while design teams iterated on these low-fidelity prototypes. While helpful, the direct participation of clinicians would better uphold the value of end-user empowerment.

We have an opportunity to re-imagine the articulation of medical making designs; Make4COVID's sewing patterns offer an alternative approach to embedding orthogonal expertise in design formats. A 3D model of a face shield does not communicate how the final product can be disinfected, whose head it will comfortably fit, or how much it will fog when worn over a surgical mask. Designers have likely considered these details, the knowledge is there, but hidden until manifested in a physical prototype. The sewing patterns, on the other hand, seamlessly articulated the sewists' orthogonal and technical expertise (e.g., the relationship between stitch density and head strap security). Craft and craftspeople present a more successful example of making with orthogonal expertise. Craft does not fit neatly into the spectrum of novice and expert makers. Rather, crafting practices can help different people express their expertise through their materials, be it experienced book binders [210], blind people with little [77] to extensive experience with textiles [53], or Make4COVID's sewist network, who collaborative designed and produced masks and headbands.

5.5 DESIGN RECOMMENDATION: DOMAIN ADAPTABLE DESIGN TOOLS

Digital fabrication is often imagined as a collaboration between the maker and machine; as a conversation with a 3D printer or sewing machine facilitated by a CAD tool. However, this framing breaks down in collaborative and interdisciplinary making, where tools facilitate collaboration among many diverse makers, not just with machines. In the context of medical making, these breakdowns inhibit safe making practices. Without engaged facilitators with interdisciplinary skills, maker communities may not be able to overcome these challenges. In the context of Make4COVID, these facilitators' efforts were extraordinary, a response to the extraordinary circumstances of the pandemic. Their work reveals opportunities for systems to step in and smooth out the connections between diverse collaborators.

Design tools for 3D printing and other forms of digital fabrication are built around the material affordances of those practices. While this approach can support careful design for these manufacturing methods, this makes it difficult for designers to articulate other critical aspects of the design that derive from orthogonal domains. Designers may recognize how features relate domain-specific objectives to physical properties of the final product: e.g., how key parameters affect fit or how wall thicknesses and materials impact disinfect-ability. However, there is no way for designers to explicitly embed this information in a model. Instead, they must articulate it externally or create a high-fidelity prototype to communicate these effects with orthogonal experts who are unfamiliar with these formats. The remainder of this thesis will demonstrate approaches to embedding these design intentions into 3D models (Chapter 7) and ways of using that design intent to solve more complex domain specific problems (Chapter 8).

5.6 conclusion

This ethnographic case study of Make4COVID is a unique example of a medical making community that highly valued safe making while including expert-amateur makers. In an effort to understand how Make4COVID managed this, we examined the formats that articulated their designs from clinicians to designers to makers. We find that these formats often separated related information, making it difficult to reliably reproduce the products. Further, the affordances of designers' preferred prototypes (e.g., sketches, 3D models) were limited when placed in the hands of clinicians, creating barriers to collaborative, iterative design. To overcome these barriers, key facilitators stepped in and relayed information between different portions of the community.

Medical making, as a form of safety-critical making in interdisciplinary settings, relies on design formats that are flexible and can express details that are critical to different makers. Clinicians need to understand how a design was made safe, how it fits to the body, and how it will be used effectively. Designers need to understand the physical properties and how the device can be produced. Makers must understand what a quality product looks like and how their individual maker practices can affect that quality. By examining how design artefacts communicate these details between diverse stakeholders, we can re-imagine design tools in ways that are more appropriate for these types of collaborative maker efforts.

6 VISION: META-TOOLS FOR CLINICAL CAD TOOLS

This marks a shift in this thesis from studies of medical making to systems that can support fabrication in these domains–the beginnings of clinical CAD tools. While these systems are motivated by the needs of medical and assistive makers, I expect the following guidelines to generalize to a wider space of domain-specific design tools. For digital fabrication to bring on its promised revolution, CAD tools need to support a wider set of users than technical experts. Supporting diverse domains with generalized tools may be infeasible; many tools will need to be domain specific.

Focusing on the design of domain specific tools raises questions about how we build tools themselves. Building CAD tools requires us to pull together complex systems from computer graphics, software engineer, and interaction design. Domain specific CAD tools will require *meta-tools* that help us embed domain expertise into existing systems and build new systems that use that expertise. The next two chapters will present two types of generalized meta-tools that can enable the construction of clinical CAD tools. In Chapter 7, PARTs enables designers to embed their design intent (i.e., their orthogonal expertise) into 3D models directly. This enables verification and reuse of 3D models by end-users who adapt these designs without the original designer's expertise. Chapter 8 presents OPTIMUM, which supports building the optimization component of generative design tools by facilitating collaboration between programmers and orthogonal experts. Chapter 9 and Chapter 10 demonstrate full applications of this toolkit to an assistive domain and a general purpose fabrication domain. The recommendations presented for clinical CAD tools provide insights into the requirements of meta-tools that will support their development, as well as development of tools for other fabrication domains.

ENABLE VERIFIABLE REUSE. Sharing and reusing digital designs is one of the primary benefits of digital fabrication. However, reuse poses risks because the changes to designs or manufacturing methods can alter critical features of the design, particularly when safety is of critical concern. The output of meta-tools is intended to be reused, but this requires mechanisms to limit unintended and preventable outcomes (e.g., broken face shields).

SUPPORT SMALL ADAPTATIONS. While the maker revolution highlights the opportunities to produce innovative products [157, 245], this is not digital fabrication's only benefit. It also enables designers to make small changes to established designs in ways that better meet end-users needs (e.g., customized splints). Generalized CAD tools can support innovation, while domain-specific tools will support small, but impactful, adaptations.

SUPPORT PRODUCTION AND PROTOTYPING. There is a valuable focus on digital fabrication, particularly 3D printing, as tools to support rapid prototyping. However, this focus should not detract from support for other stages of the design process. As medical makers consistently reiterated, prototypes are not products, and the gap between them represents a significant engineering challenge. Domain-specific design tools, and by extension meta-tools, must support the full design process and ensure that resulting products are safe and effective.

AMPLIFY ORTHOGONAL EXPERTISE. Medical and disabled makers are not novices; tools designed for novices are not sufficient. However, these makers do not necessarily have the technical expertise assumed by more advanced CAD tools; they are not often engineers. Meta-tools have two primary users, orthogonal and technical experts, and must be structured to amplify both types of expertise and foster collaboration.

7 PARAMETERIZED ABSTRACTIONS OF REUSABLE THINGS: SUPPORTING REQUIREMENTS SPECIFICATION AND REUSE

As shown by the considerations of medical makers, and the significant amounts of documentation needed to create reusable medical device designs, digital fabrication goes far beyond the geometry of a 3D model. Particularly when designs are going to be reused or adapted in new contexts, the designs need to express the designer's intentions (e.g., the way components relate or the requirements of reusing them). Modelers would benefit from the equivalent of an end user programming tool. This chapter presents the *Parameterized Abstraction for Reusable Things* (PARTs) framework, which provides a way for modelers to express their domain-expertise by building on concepts from end-user programming. PARTs supports the reuse, verification, and adaptation of 3D models used for digital fabrication.

PARTs provides the abstraction of functional geometry, which is analogous to the object oriented programming concept of classes. Functional geometry encapsulates geometric data and functionality, making it easier to validate and mutate data, manage complexity, and support modularity [55, 137]. Functional geometry supports encapsulation, validation, and mutation through two sub-types: assertions and integrators. Assertions test whether a model is used correctly, while integrators mutate designs using standard rules.

The PARTs framework is an extension of the Autodesk Fusion360 CAD tool. CAD tools already provide many helpful capabilities, including validation methods, simple geometry operations, parameters, constraints, and data structures for hierarchical composition. PARTs unites these functions under a single framework and user mental-model, making functionality easier to reproduce and use. It does this through small, but important, additions to the Fusion360 GUI that provide three important benefits.

1. Functional geometry visually represents design-intent in a 3D model rather than requiring designers to express it in separable documentations.

2. Assertions can express both design constraints and the design's intended context.

3. Functional geometry encapsulates designs in reusable and adaptable components, allowing for separation of concerns.

By leveraging PARTs' unique features, a designer can easily encapsulate information about how a design should be used. This makes direct manipulation of design intent possible, which, in turn, makes model reuse, customization, and recombination simpler and more intuitive.

7.1 SCENARIO: CREATING A CUP HOLDER MOUNT FOR A BICYCLE

An ecosystem of users can use PARTs to share designs: *designers*, who create and share complex designs, and *modelers*, who may not have technical expertise in 3D modeling, but can reuse designs for new purposes. Modelers are representative of both the expert-amateurs in hobbyist maker communities and orthogonal-experts from medical maker communities. Consider a hypothetical modeler, Kavi, who is creating a bike-mounted cup holder. Kavi views a list of the available Functional Geometry Object (FGO)s to find a cup holder and bike mount (Figure 9 left). FGOs capture design intent to make it visible to the modeler (see the semi-transparent cup at the center of Figure 9). Kavi selects the cup holder to add its geometry to the modeling space and component hierarchy. He adjusts the cup diameter to match his water bottle. Then he adds the bike clamp and positions them so they overlap. PARTs checks for violations of design intent (failed assertions) and highlights them in red during modeling. As the bike handlebar would block the cup, it is highlighted in red. Kavi changes the model until he finds a valid position, and then integrates the cup holder to join them (Figure 9 right).

The cup holder and bike mount are fairly specific FGOs that are not in the default PARTs library. Instead, PARTs supports sharing of FGOs created by designers such as Nisha. She created Kavi's cup holder FGO,

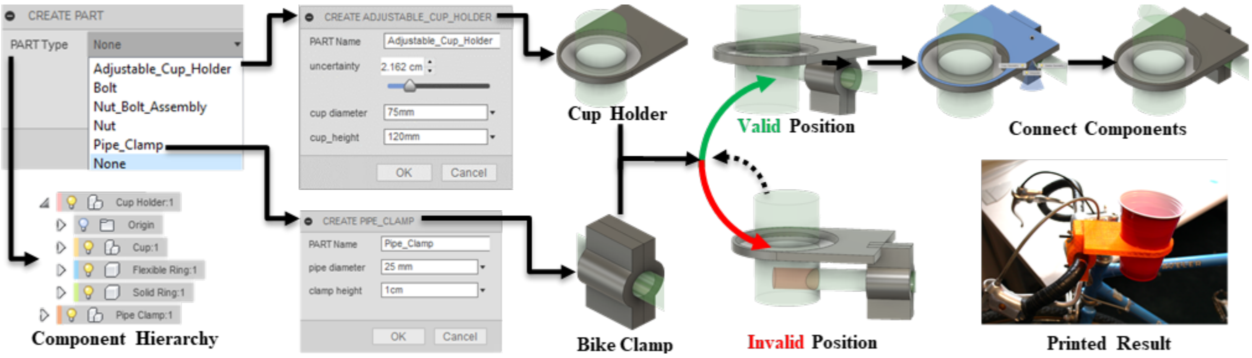

Figure 9: Modelers combine geometry and logic to define FGOs as a set of assertions and integrators. For example, to create a cup holder that attaches to a bike handlebar, the modeler instantiates cup-holder and bike-clamp FGOs from the PARTs library and iteratively modifies them in Fusion360, while the FGO assertions help to visualize potential errors. When the modeler is ready, he can integrate the cup-holder and bike-clamp to create one design.

which specifies her design intent with assertions and integrators. First, she modeled the geometry of the cup holder. Nisha specifies that this should attach to any surrounding geometry by right-clicking on the geometry and associating it with a union integrator. Next, she creates a cylinder representing the cup. She attaches the cylinder to an assertion that tests for interference, specifying that cup-holder should always leave room for the cup. It is this assertion that automatically highlights violations when Kavi positions the FGO.

Suppose a different designer, Laura, created Kavi's handlebar mount. Laura models the clamp and adds an assertion that ensures space for the handle bar. She also adds a fastener to close the clamp. The PARTs library includes a fastener FGO with bolt length and diameter parameters. The fastener FGO ensures there is enough material around the bolt, and that the bolt has a clear path. During integration, this FGO will create a bolt hole. All of this is now part of the handlebar FGO and will help Kavi when he combines the bike-mount with the cup holder.

This scenario highlights how the PARTs framework engages multiple types of users in a synergistic way. Similar to the user types proposed in [162], PARTs supports the transfer of designs between programmers who extend the framework, designers (Nisha, Laura), and modelers (Kavi). It also demonstrates how PARTs supports reuse at multiple levels, and enables the creation of complex, hierarchical objects.

7.2 THE PARTS FRAMEWORK

The PARTs framework allows models to include design intent, making them easier to reuse. Through small changes to the Fusion360 GUI, designers and modelers create functional geometry. The PARTs framework is implemented using Fusion360's API to leverage standard capabilities for creating and manipulating geometry, such as component hierarchy and interference testing. Because PARTs is deeply integrated with Fusion360, using PARTs involves a combination of Fusion360 and PARTs commands embedded in the GUI.

7.2.1 *Functional Geometry*

Functional geometry visually captures design intent so that the modeler can avoid problems as models are combined. From a modeler's perspective, functional geometry is analogous to a Fusion360 component. Fusion360 organizes the modelers design by components. Just like components, functional geometry makes use of geometric parameters and features, but goes further than standard parameters and constraints to visualize violations of assertions that represent semantic expectations about models (see Figure 9 center).

A critical feature of PARTs is that a designer uses FGOs to express semantic concepts geometrically and without programming. An example is adding a cup model to a cup holder. When shared with a modeler, this reveals important considerations: 1) do not interfere with the cup, and 2) size the hole to match the cup.

The ability to use standard operations to document design intent is the feature that makes PARTs unique. It had a major impact on ease of use in the user study workshops, even though it does not automatically resolve violations.

From a programmer's perspective, an FGO is a data structure that inherits from the PARTs base FGO class hierarchy. It includes code for generating geometry and for operating on that geometry based on design intent. Each FGO includes two special types of components. Assertions check that the FGO is valid, not violating some part of the design intent. Unlike constraints or parameters, assertions do not enforce a set of rules. Instead, they visually highlight violations. For example, an assertion might check that the model is not intersected by other geometry. Integrators mutate the model, such as cutting a hole for the FGO or generating a connection between the FGO and the surrounding model. Assertions and integrators can act on standard model geometry as well as other FGOs. Assertions and integrators are complementary, in that assertions check and report on design intent violations, while integrators enforce design intent.

Modelers instantiate a new FGO by opening a dialog showing the PARTs library. At instantiation, the default elements' geometry is generated and added to the component hierarchy. A dialogue is shown if parameters are needed to generate geometry used by the FGO's assertions and integrators. The FGO's geometry can also be customized, and the modeler can add new assertions and integrators. A designer can create a novel FGO without programming. The designer starts with an FGO containing no assertions or integrators, and progressively adds them from the library.

Internally, PARTs manages a symbol table that associates the component and FGO. When the modeler right-clicks on a component, PARTs uses this table to find the associated FGO. If an FGO is found, the menu displays commands to copy, delete, add functionality to, and integrate the FGO.

7.2.2 *Assertions*

The purpose of assertions is to test whether a design meets the original designer's expectations, even when a modeler alters the model. For example, an assertion might test for space for an object, ensure there is a path for tools needed in assembly, or check for the presence of a related part. Each assertion takes a geometric parameter as input and runs a Boolean function against that parameter. PARTs includes two subclasses, which test interference and overlap of the geometric parameter against the surrounding geometry. *Interference* fails if there is an intersection, such as when a path for a tool is blocked. *Overlap* fails if any part of a geometric parameter is not intersected. For example, a privacy cover for a camera might associate an overlap assertion with the portion of the model in front of the camera lens.

The PARTs framework automatically tracks changes to geometry and checks any assertions that intersect the changed geometry. Failed assertions are highlighted in red. PARTs listens for edit events in the Fusion360 UI which signal model changes. Assertions that intersect modified geometry are tested after each change, since this is where assertion violations will be introduced.

PARTs does not impose solutions, interrupt modelers, prevent errors, or optimize designs to solve problems. Instead, PARTs provides visual information about design intent. PARTs works in concert with Fusion360. Stronger design requirements can be captured using Fusion360 tools, such as parameters or constraints. PARTs' complementary approach allows non-experts to understand and manipulate design intent. By not strictly enforcing assertions, users are allowed to experiment with their design intent, something participants in the workshops praised.

7.2.3 *Integrators*

While assertions test design expectations, integrators use geometry and other parameters to enact design intent. Integrators can be combined so that an FGO might include one integrator that puts a mold around a phone and another that attaches the phone holder to the surrounding geometry. Each integrator takes a geometric parameter as input and runs a function that modifies the surrounding model against that parameter. PARTs provides two subclasses: *union* and *cut*. By combining multiple integrators, FGOs can execute a wide range of mutations to integrate with the surrounding model.

Integrators are defined by the designer of an FGO. The integration geometry helps to indicate design intent. Modelers activate integrators by right clicking on an FGO and selecting the integrate command.

This converts the integration geometry into a feature of the surrounding model (e.g., cutting a hole for a bolt). Although integration mutates the model, Fusion360 has sophisticated support for undo, which PARTs leverages. A modeler can delete/edit the integration actions, which are labeled in the Fusion360 timeline, or re-invoke integration.

7.3 LIBRARY OF REUSABLE DESIGN PATTERNS

We provide a small library of assertions and integrators that can be combined and recombined to express a variety of reusable 3D modeling design patterns. We derives these design patterns from a survey of 3D printable designs found on the Thingiverse 3D model repository.

7.3.1 *Survey Method*

We surveyed 10,560 designs on Thingiverse, gathered in two phases in 2016 and 2017. Each year, we gathered 40 pages of 12 Things from Thingiverse's 11 categories (5,280). 20 pages of results were the most popular designs at the time of search, and the other 20 were the newest designs. Of these, we kept the 962 designs that met the following requirements:

- Success: Fabricated, or marked as 'made' by a different Thingiverse user (8,839, 83.70%).
- 3D Printable: Intended for 3D printing (9,257, 87.66%).
- Non-Trinket: Images and descriptions show interaction with existing objects. (1,246, 11.81%)

Using affinity diagramming [112], we identified design patterns for incorporating existing objects into models. Patterns demonstrated in Figure 10 implicitly express designer's intentions.

(a) Approximate (682 models)

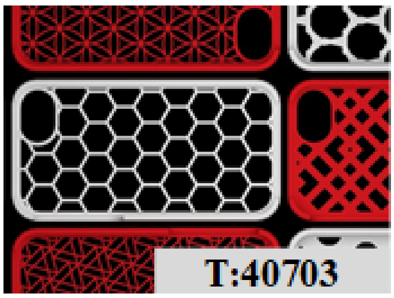

(b) Mold (219 models)

(c) Hold (51 models)

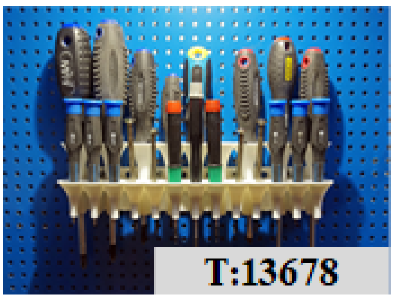

(d) Connect Similar Objects (225 models)

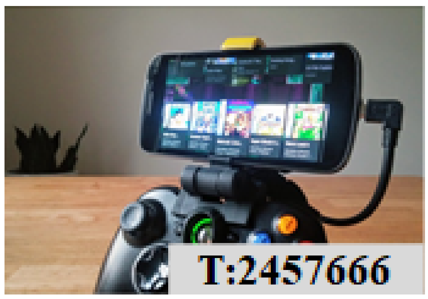

(e) Connect Different Objects (159 models)

Figure 10: Example Thingiverse Designs for each design pattern. Labels show 'Thing' number used to identify designs in the repository.

7.3.2 *Design Patterns*

Table 5: Design patterns used to create sample objects with FGOs from the PARTs library.

	Survey Example	PARTs Model	Printed Result
Bounding Box			
Scaled Mold	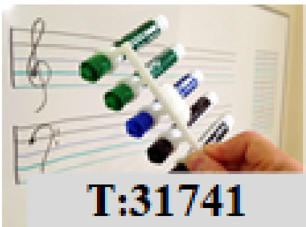		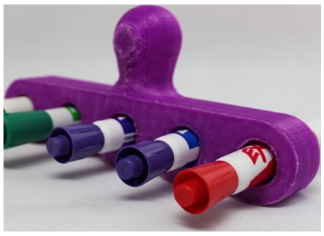
Swept-Path			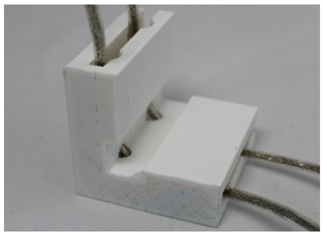
Connector	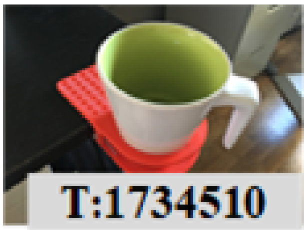		

Based on the salient patterns from the survey, we developed four FGOs to express these patterns with the PARTs library. In three cases, the design pattern is supported by an integrator that cuts a hole for an existing object, with two related assertions that test that the hole remains empty, and is surrounded by a minimum amount of material. These inherit from the cut integrator, interference assertion, and overlap assertion, respectively. The fourth case, connect, uses a single integrator, which generates material connecting two faces.

While the design concepts behind these library elements could, in model-specific ways, be implemented without PARTs, it is PARTs that imbues them with both usability and re-usability through encapsulation. These powerful library elements can express a wide range of design intentions.

BOUNDING BOX: Many designs in the survey approximate real-world objects using rectangular holes. To support this, the box-cut integrator, interfered-box assertion, and minimum-boundary box assertion each generates an object's bounding box around a specified object. The integrator cuts out space for this bounding box, and the assertions test that the space is not interfered and that there is a compliance buffer around it. The first row of Table 5 shows a thumb drive/smartphone holder from the survey, (left), the PARTs model (center), and the printed result (right).

SCALED MODEL: The scaled-hole integrator, interfered-model assertion, and minimum-surrounding material assertion support a mold design pattern. The scaled-hole integrator cuts space for an object and the interfered-model assertion ensures it stays clear. Each takes a model of the real-world object as a parameter. The minimum-surrounding material assertion checks that material surrounds the hole. This then takes the same geometry as input, but scales it to a thickness defined by the user. The second row of Table 5 shows a drawing tool modeled with these elements.

SWEPT-PATH: Some designs, such as the cable organizer in the third row of Table 5, use cut-paths to hold flexible objects. The path-cut integrator, interfered-path assertion, and minimum-boundary around path assertion reserve and create paths for objects along a user-defined curve.

CONNECTION INTEGRATOR: The remaining design patterns connect two objects. The Connection Integrator is parameterized by a starting and ending face, and grows a structure between them using an adapted Steiner Tree support-generation algorithm [254]. The last row of Table 5 shows the support structure holding a hefty mug full of coffee.

7.4 DEMONSTRATIONS

The following demonstrations show how complex FGOs, created without programming, support a variety of design-reuse activities.

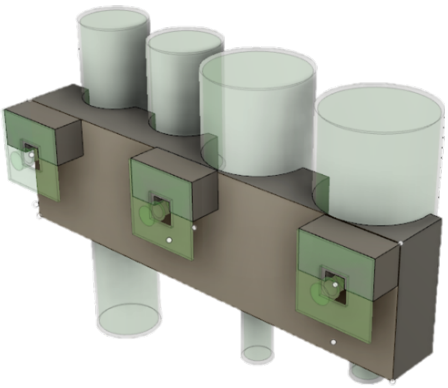

(a) Model of pegboard organizer

(b) Printed result of pegboard organizer

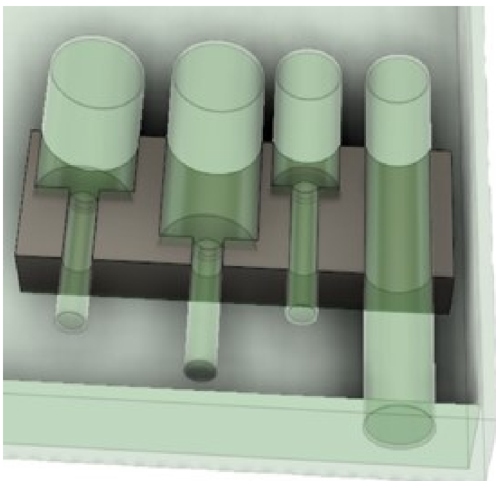

(c) Model of organizer in a drawer assertion

(d) Printed result of drawer organizer

Figure 11: Small changes between models support adaptation of the organizer to a new context.

7.4.1 *Composing Functional Geometry to Create a Tool Organizer*

In this demonstration, Nisha designs a tool organizer, and uses functional geometry to express her design intent. This shows how functional geometry can help designers separate their design into sub-tasks (e.g., create a pegboard hook; use a fastener, organize tools). These sub-components make the design reusable and adaptable.

This example starts with a high-level goal: organizing tools by hanging them from a pegboard. Nisha breaks this problem into a series of simpler tasks: build a pegboard hook, incorporate a bolt and nut into the hook, build an organizer, add hooks to the organizer, organize tools in the organizer, add multiple tools to the organizer. The entire design only requires basic modeling skills, creating boxes and cylinders.

Nisha first defines a pegboard-hook FGO with a union integrator parameterized by a small box to cover a peg hole. She also defines one assertion that ensures nothing in the model interferes with a thin box, representing the pegboard. Next, Nisha adds a fastener FGO to the pegboard-hook. She positions the nut inside the box, while the head of the bolt protrudes. The bolt head will function like a peg (Figure 11a). At integration time, integrators associated with the fastener FGO will cut holes for the bolt and nut. Having created a hook, Nisha solves the larger problem of using many hooks to hang an object from the pegboard. She creates a box to hold tools and makes copies of the hook on the back of the box.

Next, Nisha defines a tool FGO, using a Scaled-Hole integrator and interference assertion parameterized by an approximation of the tool's shape. She makes slightly altered copies of this FGO to represent different tools. Then, she positions the tool FGOs in the organizer box. When they are instantiated, tool-shaped holes will be created.

Suppose that the modeler, Kavi, wishes to reuse the tool organizer to fit into a drawer, rather than hang from a pegboard (Figure 11d). He first creates a hollow box, representing the drawer, with an interference assertion. This ensures that the organizer fits completely within the drawer. He places the organizer inside the box, which triggers an error highlighting the pegboard hooks that intersect the drawer. This prompts him to remove the unnecessary hooks. With these two small changes, Kavi has adapted the design to reuse it in a new context.

7.4.2 *Expressing Design Intent in a Complex Lamp Assembly*

PARTs is powerful enough to describe complex objects. To illustrate this, we recreated an example from the Thingiverse survey, a lamp composed of CDs, bike spokes, a bulb, and 3D printed components. Although this design involves many real-world objects, the Thingiverse model includes just the 3D printed geometry. In contrast, PARTs describes the whole design, including the use of the non-printed objects. This makes the final model easier to understand and reuse (Figure 12).

The designer, Laura, first creates FGOs for a CD and bike spoke. She uses the scaled-hole integrator and a matching interference assertion to reserve space for the spoke. Next, Laura creates a new FGO, a connector, which attaches the CD to the bike spoke. To this, she adds the spoke FGO and a fastener FGO. She customizes the fastener FGO by adding a washer to connect the cylinder to the CD. These FGOs will cut holes in the connector, creating an assembly of the spoke and CD.

Next, Laura creates a hub FGO, with an arm for each spoke. To this hub, she adds an interference assertion to reserve space for a bulb socket, and adds a swept-path interference assertion and the corollary integrator to represent the cable. Laura creates and adds a light-bulb FGO. The bulb FGO uses an interference assertion and a sphere around the bulb to ensure that no meltable plastic is too close to the hot bulb.

This example illustrates how functional geometry not only improves modularity and specifies components for 3D printing (the connectors and central hub), but also shows how the components inter-operate with real-world objects (CDs, bike spokes, nuts, bolts, washers, light-bulb, socket, and cable). The FGO hierarchy created by the designer, Laura, lays out the relationships between the various objects, making it easy to alter the design while maintaining her design intent.

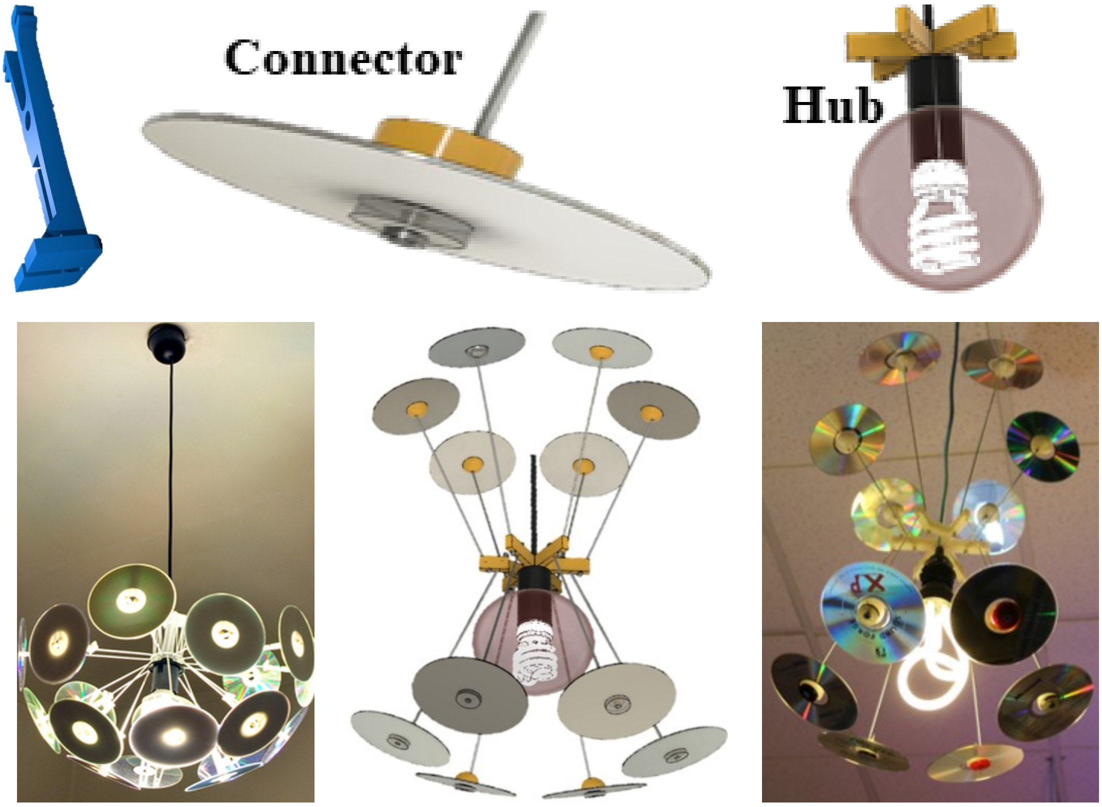

(d) Original Design (e) Full PARTs Model (f) Full PARTs Photo

Figure 12: Functional geometry contextualizes the modeled components among existing objects.

7.4.3 *Extensibility of PARTs through Scripting*

Many CAD tools can already be extended with scripted plugins; indeed, PARTs extends Fusion360. PARTs supports the encapsulation of new features by programmers, because it provides a simple framework which enables programmers to create new library elements that express general concepts, are accessible through a unified interface, and can be shared widely. In addition to making 3D models easier to reuse, PARTs is a meta-tool that helps programmers to integrate specialized 3D model components (i.e., FGOs) into a generalized CAD tool. Here, we illustrate how PARTs enables extension of Fusion360 with a fastener wizard and an uncertainty buffer FGO generator.

FASTENER WIZARD Fastener Wizards are common in CAD programs. For example, the Solidworks "Smart Fasteners" [2] tool inserts representations of fasteners into an assembly and validates the configuration. An advantage of PARTs is its ability to validate and integrate the fastener in one modeling stage.

The fastener FGO is specific to nuts and bolts. It consists of FGOs for the nut and for the bolt. Each child FGO has an interference assertion that reserves space for the hardware, a cut integrator that cuts holes with surrounding clearance for the hardware to slip into, and an additional overlap assertion that checks for material surrounding the resulting holes. By sharing a parent FGO, the bolt and nut FGO are able to share parameters (bolt diameter and length) and be moved through the design as one component, maintaining their respective alignment.

Like a conventional fastener wizard, the fastener FGO can integrate itself into the surrounding model by cutting an appropriately sized hole. Unlike traditional tools, the fastener FGO can be extended to express new design intents, such as a test that reserves space above the bolt to ensure a tool can reach it during assembly.

UNCERTAINTY BUFFER GENERATOR Moving beyond fasteners, PARTs helps us extend Fusion360 with an uncertainty buffer generator, based on Kim *et al's* technique for managing uncertain object measurements[133].

The script generates a buffer FGO that approximates a real-world object, given a measure of uncertainty around that object. For example, a modeler could measure a cup three times and enter the diameter range. The FGO uses two integrators to create this mechanism. The first integrator generates a ring around the model that has a radius of the uncertainty parameter. The ring remains separate from the model and would be printed with flexible materials. The second FGO uses a union integrator to generate a hard ring that surrounds the buffer, so the FGO creates a cup holder that accommodates mugs of different sizes and shapes.

7.5 EVALUATION: MODELING WITH DESIGN INTENT

To validate that the PARTs framework would allow non-expert modelers to create reusable and validatable designs, we held two workshops with a total of ten participants (Table 6) who each had previous CAD experience. The first workshop was conducted with six members of a makerspace (ID's starting with A). The second workshop was conducted in a research lab with 4 graduate students (ID's starting with B). Prior to attending the workshop, all participants were given links to Fusion360's tutorials, with notes on which tutorials would be most valuable for the workshop. 6 of the 10 participants used these tutorials, and 2 participants had pre-existing experience with Fusion360. The remaining participants, A3 and B3, had no Fusion360 experience prior to attending the workshops, but they had more CAD experience than most other participants.

Table 6: Participant Demographics

ID	Gender	Age	Years of Experience	CAD Tools	Profession
A1	Male	36	3	Other	Software Engineer
A2	Male	31	10	Fusion360, SolidWorks	Graphic Artist
A3	Other	40	20	Blender, OpenSCAD, SolidWorks	Startup Owner
A4	Female	30	1	Blender	Cosplay Designer
A5	Male	33	1	Blender	Defense Contractor
A6	Male	28	4	OpenSCAD, SolidWorks	Software Engineer
B1	Male	28	1	Fusion360, SolidWorks	Design Student
B2	Other	23	2	Blender, SolidWorks	CS Student
B3	Male	26	8	SolidWorks	CS Student
B4	Male	33	2	OpenSCAD	CS Student

7.5.1 *Method*

The workshops consisted of four phases: discussion (20 minutes), training (20 minutes), think-aloud modeling (60 minutes), and review (20 minutes). The sessions were audio recorded and photographed for further review.

During the training phase, we demonstrated the interface features of PARTs and Fusion360 by running through the simple example of joining two blocks with a bolt and nut. The demonstration included: using an FGO from the library, checking assertions, and invoking integration. We also provided descriptions of the other FGO's available in the library. We then reviewed how to complete the same task without PARTs.

During the modeling phase, participants were given a modeling goal, example objects, and measurement tools. Participants were allowed to speak to one another and ask researchers for assistance with the

Fusion360 and PARTs UI or for clarification on the design task. Researchers did not give advice on how participants could complete the task, but merely supported using the tools.

During the review stage, participants gathered into a group to discuss their experience. They were prompted to discuss how the PARTs framework differed from their previous experience with CAD tools, what criticisms of the tool they had, and if it affected their modeling behavior. Participants in the later workshops were also asked to compare their experiences using FGOs versus standard geometry.

The first workshop tested designers experience's of PARTs. Participants were tasked with creating a cup-holder using PARTs. Modelers had access to all assertions and integrators, but no pre-existing FGOs to reuse in their design. Successful completion of the task required that the resulting model fit the provided cup, and that the holder was stable when the cup was filled.

The second workshop focused on modelers, comparing PARTs to Fusion360. Participants were asked to reuse models to create cup-holders that could be mounted on a bike. They were placed in two conditions, Fusion360 (alone) or PARTs (Fusion360 with PARTs). Condition ordering was randomly assigned to ensure that participants were not primed by the first task to perform better in the second. For the Fusion360 condition, we provided parameterized models based off a cup-holder (T:2271973) and bike-mount (T:1012573). In the PARTs condition, we provided a parameterized cup-holder and bike-mount FGO. To complete the task, participants needed to: correctly parameterize the cup-holder, bike-mount, and fastener; position the cup-holder and bike-mount components; position and cut space for the fastener.

7.5.2 *Findings*

Anecdotes from participants in both studies show the benefits of PARTs. For example, A1 noted that he often modeled real world objects to help with his modeling process. His process was manual and informal, and he appreciated the automated support for capturing these process ideas, saying

> *"It needs to be in your model for the design or simulation to mean anything. " (A1)*

Participants in the second workshop voiced frustration about Fusion360, precisely because it lacked these descriptions. All participants had questions about how the cup and pipe fit into the model, but no similar questions arose while using PARTs. B3 complained that the cup-holder model was under-specified.

> *"I guess that's where the cup goes, I'm not sure how to fit it to this cup " (B3)*

The ability to delay integration was also valued, in comparison to a standard process where

> *"...you really have to know what you intend to build everything right " (A6)*

A2 was excited by the idea that he could use integrators as place holders for changes he wanted to make later, just as he uses layers in photo editing to prepare changes to the final image. Again, participants in the second workshop noted the lack of these capabilities in Fusion360.

> *"It's essentially a 3D version of layers. It's like a non-destructible field on one side and a visual debugger on the other " (A2)*

7.5.2.1 *Workshop 1*

4 of 6 participants completed the PARTs design task in the first workshop, and 3 of those models were completely stable. All participants used similar design patterns to create their model; each represented their cup as an interference assertion. They also created cut integrators with a copy of the same geometry. The three failure cases diverged when creating a base for the cup holder. Participants A3 and A6 ran out of time while making an accurate cup model. A4 completed the design task but did not accurately measure the cup, causing a slightly unstable print. Images of the resulting cup holders are shown in Table 7, Row 1.

Table 7: Workshop model and print results. Designs that could not be printed are not shown

	A1	A2	A4	A5
W1: Prints				
	B1	**B2**	**B3**	**B4**
W2: F360				
W2: PARTs				
W2: Prints			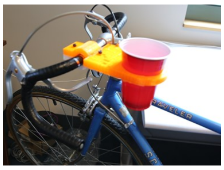	

Table 8: Participant success and failure in modeling tasks.

Participant	Condition	Fits Cup	Fits Bolt	Fits Bike	Component Placement	Fastener Placement	Printable and Functional
B1	PARTs	Pass	Pass	Pass	Pass	Fail	Fail
	Fusion 360	Fail	Fail	Fail	Pass	Fail	Fail
B2	PARTs	Pass	Pass	Pass	Pass	Pass	Pass
	Fusion 360	Fail	Fail	Pass	Pass	Pass	Fail
B3	PARTs	Pass	Pass	Pass	Pass	Pass	Pass
	Fusion 360	Pass	Fail	Fail	Fail	Fail	Fail
B4	PARTs	Pass	Pass	Pass	Pass	Fail	Fail
	Fusion 360	Pass	Fail	Fail	Fail	Fail	Fail

7.5.2.2 *Workshop 2*

Participant failure rates are shown in Table 8 broken up over six sub-tasks. Under the Fusion360 condition, no participant created a printable model, and the participants averaged completion of 1.5/6 steps. Under the PARTs condition, B2 and B3 were able to create printable and functional models, and participants averaged 5/6 completed steps. Within subjects, they improved by an average of 3.5 steps.

The main source of failure for participants under the Fusion360 condition was adjusting the parameters of the pre-existing models to fit the new objects. This was a challenge for two reasons. First, despite being shown parameter definitions in the tutorials, participants could not figure out how to change the dimensions of the models, and instead used a scaling feature to adjust the model size as a whole. Second, participants did not think about other effects on the parameters besides the objects dimensions, such as leaving space for the bolt. Neither of these were challenges for the PARTs condition, where all participants succeeded in all parameterization tasks. This is likely because FGOs afford clear parameter declarations that are displayed at instantiation.

B3 and B4 were further challenged by attaching the cup holder and handlebar mount, and ran out of time trying to join the two models together. Compounding errors and a rush to complete the tasks in time resulted in B3 and B4 not attempting to place their bolts, and B1 placing the bolt in an infeasible position. These challenges were not completely surmounted with PARTs, but all participants were able to position and combine the cup-holder and bike-mount FGOs. This is likely because they did not feel rushed, as the parameterization task were significantly easier, though B1 and B4 still found it difficult to position the fastener. Under both conditions, these participants struggled with Fusion360's move command, leading to failures to position models.

7.5.3 *Discussion and Limitations*

Participants found the following features of PARTs most accessible: adjusting parameters, integrating models, and making functional changes to another modeler's design. While PARTs helped participants complete their task more effectively, it was not without frustrations. Participants complained that assertions were only occasionally helpful and wanted to turn the visualizations off.

There were some limitations to the second workshop that may affect the results. For instance, participants struggled to contextualize the Fusion360 models, which may not have been true if they had discovered them online, as they would in real life. Further, users with more Fusion360 experience may have shown less of a difference between conditions.

7.6 CONCLUSION

Most work on CAD tools has focused on helping designers express the shape of what they will fabricate, with little regard for the intentions behind those shapes or how they will be modified by future users. Small changes made by an uninformed modeler can cause critical failures in how the design is used. Designers may take the time to document their intentions, but all too often, this documentation is lost when a model is being adapted. Functional Geometry relates geometry to simple programs that modify the modeling environment. This can support testing and modeling, and gives designers new ways to express their intentions directly in the 3D model.

Beyond making individual designs more reusable, PARTs creates a new way of extending CAD tools. Programmers can provide scripts that operate on geometry as though it is a function-parameter. Doing so makes it easier to represent algorithmic model changes without building a tool extension from scratch. This serves as a flexible jumping off point for defining classes of reusable designs, rather than single reusable models.

8 OPTIMUM: DOMAIN-SPECIFIC GENERATIVE DESIGN WITH METAHEURISTIC METHODS

While PARTs focuses on making designs more reusable, the challenge of creating designs from scratch is still significant for many medical makers. Designing medical and assistive devices from scratch is a difficult task that requires teams of technical and orthogonal experts. For instance, standardized splint templates are the result of research and practice from a wide variety of experts (e.g., occupational therapists, rehabilitation engineers). Even amongst standardized designs, medical makers face challenges when selecting or customizing devices for end-users' needs. Rarely does the process of selecting solutions among standardized designs align with the affordances of CAD tools that are designed to support novel invention.

Instead, medical makers need domain-specific tools that explore the space of available designs. Generative design can help designers create unique designs by specifying goals (e.g., fit), rather than directly creating the design. This is a powerful approach that can help designers rapidly explore a design space and, potentially, automatically realize their goals by amplifying the expertise of domain and programming experts. Building these tools requires the expertise of programmers who tailor optimization methods to the domain, and orthogonal experts who provide the insights necessary to make the optimization effective. This complex collaboration between programmers and orthogonal experts is critical, but largely undersupported. Existing systems tend to only support the programmer, by offering off-the-shelf implementations of standard optimization algorithms which need to be tailored to a domain. However, the process of incorporating orthogonal expertise is left entirely up to the programmer. This leaves open opportunities to support the implementation of domain specific optimizers by enabling collaboration between orthogonal experts and programmers.

The *Optimization Programming Toolkit Integrating Metaheuristic User-driven Methods* (OPTIMUM) helps two types of users, *orthogonal experts* and *programmers*, to collaboratively implement the optimization component of generative design tools in diverse domains. The resulting *optimizers* are used by a third type of user, *designers*. These designers, like orthogonal experts, use their domain-specific expertise to solve problems with the optimizers, but do not do not have the collaborative support of a programmer, and use optimizers entirely through a simple, automatically generated user interface. We have developed OPTIMUM with three principles in mind. First, OPTIMUM empowers orthogonal experts to participate in the implementation of optimizers. Second, it helps programmers and orthogonal experts to rapidly prototype a satisfactory and sufficient optimizer. Third, OPTIMUM amplifies the expertise of programmers and orthogonal experts by producing domain-specific optimizers that are accessible to designers.

OPTIMUM is a domain-agnostic toolkit that deconstructs metaheuristic optimization into a small set of useful, pluggable operations. Implementing an optimizer with OPTIMUM requires programmers and orthogonal experts to implement simple, modular, domain-specific functions, called *objectives* and *modifiers*. These components help orthogonal experts express their goals, as well as modification strategies that can help meet those goals. OPTIMUM provides a domain-agnostic library of pluggable components, that help programmers rapidly prototype domain specific optimizers to apply these objectives and modifiers. OPTIMUM generates a simple GUI for each optimizer, which designers use to generate unique designs.

OPTIMUM uniquely enables programmers and orthogonal experts to collaboratively implement the critical optimization component of generative design tools. Unlike existing optimization toolkits (e.g., [83, 166, 194–196, 201, 249]), OPTIMUM assists programmers and orthogonal experts in tailoring methods to their domain. Further, unlike most generative design tools (e.g., [46, 72, 75, 168, 226]), OPTIMUM's optimizers can be readily modified by designers (i.e., end-users) without the assistance of a programmer.

In this paper, we describe how OPTIMUM implements metaheuristic optimizers and how can be used by orthogonal experts, programmers, and designers. OPTIMUM is particularly targeted at generative design tasks that tune standardized designs to meet a specific person's needs. In this chapter, we demonstrate three tools to show how OPTIMUM; (1) supported an Ophthalmologist and programmer in collaboratively building a cataract lens selection tool, (2) enabled us to derive a thumb splint optimizer from occupational therapists' descriptions of the splint design process that tailors splints to a patient's needs and preferences, and (3) enabled us to rapidly prototype different metaheuristic optimizers that replicate an existing generative design tool [44] and make it more efficient and effective. Chapter 9 and Chapter 10 will build on OPTIMUM with more extensive generative design tools and evaluations of their use.

8.1 OPTIMUM TOOLKIT

OPTIMUM is a toolkit that helps orthogonal experts and programmers collaboratively implement domain specific optimizers, which can be used and modified by designers, without programming, through a simple user interface. It is built around domain-specific heuristic libraries that are developed by the orthogonal experts and programmers to express design strategies in the domain. These heuristics are iteratively applied by metaheuristic optimizers, which programmers configure with elements of OPTIMUM's domain-agnostic library. In the following sections, we describe the three guiding principles behind OPTIMUM and its various components.

Throughout these sections, we will refer back to the following scenario to provide concrete examples of how OPTIMUM is constructed and used. Consider two characters: the orthogonal expert, Chef Alton, who helps the programmer, Steve, to build our cookie optimizer (Figure 13). We derive this example from the transcript [177] of a cooking show that describes how to customize cookie recipes to have different textures (e.g., crispy, chewy, cakey). While this is a rather trivial and simple domain, it allows us to concretely illustrate the role of each of OPTIMUM's users and components.

Figure 13: Cookies optimized to different textures.

8.1.1 *Guiding Principles*

The tools OPTIMUM helps implement are targeted at three clinical CAD requirements: small adaptations, support for production, and amplification of orthogonal expertise. To do this, OPTIMUM's guiding principles focus on including orthogonal experts in the design of the tools wherever possible.

Empowerment of orthogonal experts: OPTIMUM empowers orthogonal experts to help create generative design tools without programming. This is why we focus on embedding orthogonal expert's heuristics directly into the optimizer, and why we have made it possible for non-programmers (i.e., orthogonal experts, designers) to modify aspects of the optimizer through a GUI.

Flexible Prototyping of Optimizers: Different metaheuristic strategies significantly affect the quality of optimizers. Which strategy is best is hard to predict and depends on the domain. To help orthogonal experts and programmers find an effective and satisfactory optimizer, we have made OPTIMUM's pluggable optimizer structure capable of expressing a wide variety of metaheuristic optimization solutions. This flexibility makes it easy to iterate on and test different solutions; trial and error comes at a very low cost.

Expertise Amplification: Design tools amplify the technical and orthogonal expertise of the teams that implement them. Through these tools, that expertise is accessible to designers (e.g., medical makers) who cannot program. OPTIMUM helps programmers and orthogonal experts collaboratively build generative design tools and, by extension, provide their expertise to the designers that use those tools.

8.1.2 *System Overview and Definitions*

Algorithm 1 The iterative structure of an OPTIMUM optimizer.

Input $seeds$: a set of starting designs provided by designers
Input $\mathbb{H}$: The heuristic map set by domain-experts and/or designers
Output $\mathbb{D}$: the population of generated designs

$\mathbb{D} \leftarrow \{\}$
for $d \in seeds$ **do** ▷ Evaluate Seed Designs
  $s_d \leftarrow evaluate(d)$
  add d to $\mathbb{D}$ sorted by s_d
$STOP \leftarrow False$
while not $STOP$ **do** ▷ Main Optimization Loop
  $d \leftarrow select\_design(\mathbb{D})$
  $m \leftarrow select\_modifier(d, \mathbb{H}, \mathbb{D})$
  $d' \leftarrow m(d)$
  $s_{d'} \leftarrow evaluate(d')$
  add d' to $\mathbb{D}$ sorted by $s_{d'}$
  $STOP \leftarrow stop(\mathbb{D})$
return $\mathbb{D}$

Optimizers consist of pluggable components (Table 9) from a domain-specific heuristic library and a domain-agnostic library. They follow a five step iterative process (Algorithm 1). The designer inputs a small set of seed *design representations* into the optimizer. First, these designs are scored by an *objective function* which orthogonal experts and designers construct from *objectives* from the heuristic library. The evaluated designs are added to a *design population*, which has a limited capacity set by the programmer. As that capacity is exceeded, poor performing designs are removed. Second, a domain-agnostic *design selector* chooses a design from the population to be considered for the next iteration. Then a domain-agnostic *modifier selector* chooses a domain-specific *modifier* to apply to the selected design. The modifier changes the design, producing a new design, which is evaluated and added to the population. Next, the optimizer considers if the population of designs meets a *stopping criteria*. If so, the population is returned to the designer. Otherwise, the cycle repeats. The following sections describe the components of an optimizer, organized by the library they derive from.

8.2 DOMAIN SPECIFIC HEURISTIC LIBRARY

Orthogonal experts and programmers implement the domain-specific components of their optimizer and organize them in a heuristic library. The orthogonal expert first describes how to represent designs, and the programmer creates a corresponding data structure called a *design representation*. OPTIMUM does not require a specific structure for design representations, since they are only used by objectives and modifiers, which are implemented by the programmer. For example, Chef Alton describes cookies as recipes and Steve implements a Python class with parameters for each ingredient amount. Then, the orthogonal expert describes design goals and how to evaluate them. When making crispy cookies, Chef Alton uses butter instead of shortening because butter has "a relatively low melting temperature so the batter spreads before setting" [177]. Steve implements this as a *low melting-point* objective, which evaluates to 1 if butter is used. The orthogonal expert also describes how they modify designs. For example, Alton switches between butter and shortening, so Steve implements a modifier that switches the type of fat used. Given this heuristic library, orthogonal experts relate objectives to modifiers in a heuristic map to express design strategies.

Table 9: A summary of components of an OPTIMUM optimizer.

Term	Definition	Notation	Example
Domain-Specific Heuristic Library	Library of domain-specific elements implemented by orthogonal experts and programmers.		
Design Representation	A data structure that represents designs in a domain	d	A set of parameters that define a cookie recipe
Objective	A function that evaluates how well a design meets a specified criteria or design goal	$0 \leqslant o(d) \leqslant 1 \forall o \in \mathbb{O}$	A function that compares a cookie recipe's estimated melting point to a target value set by a designer.
Objective Function	A function that evaluates the weighted set of objectives to assess a design's quality	$evaluate(d)$	see Equation 1a
Modifier	A function that creates a new design by modifying a generated design.	$m(d) = d' \forall m \in \mathbb{M}$	A function that sets the fat used in a recipe to butter
Heuristic Map	A set of weights between modifiers and objectives that express how effective a modifier is expected to be at improving an objective	$\alpha_{m \to o} \in \mathbb{H}$	A designer sets a weight of 2.0 between a low-melting point objective and a butter modifier
Domain Agnostic Metaheuristic Library	Library of domain agnostic components of a metaheuristic optimizer		
Design Population	A data structure that organizes generated designs	$\mathbb{D}$	see Section 8.3.1
Design Selector	A function that selects a design from the design population	$select\_design(\mathbb{D}) = d$	A function which returns the highest scoring design that has been generated.
Modifier Selector	A function that selects a modifier to use on a design	$select\_modifier(d, \mathbb{H}, \mathbb{D}) = m$	A function that selects the modifier that is expected to most improve the design
Stopping Criteria	A function that determines if the optimization results should be returned	$stop(\mathbb{D}) = \{True, False\}$	A function that returns True if the design scores more a threshold value.

8.2.1 *Objectives*

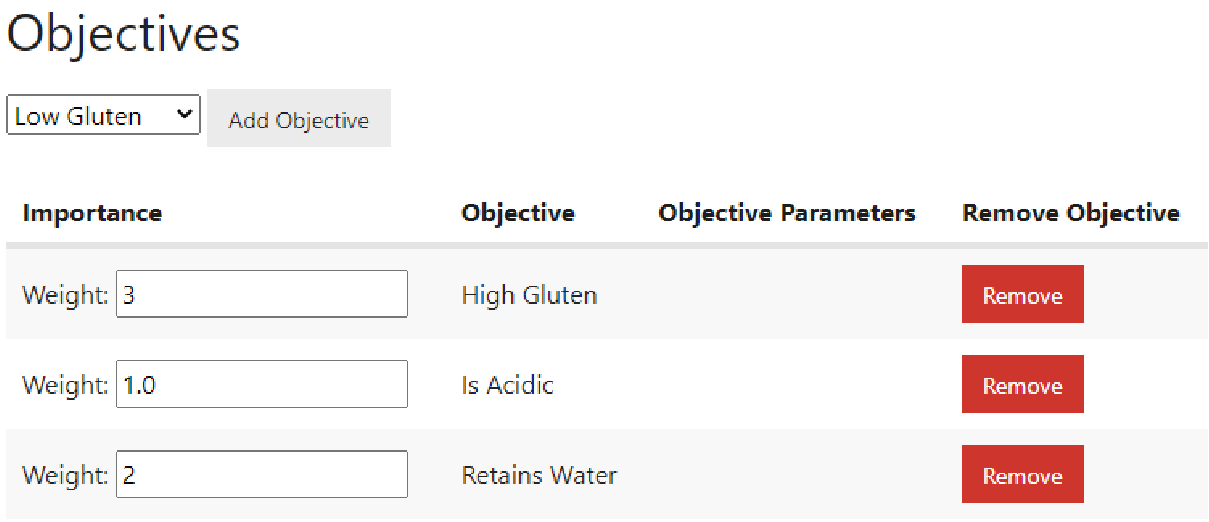

Figure 14: Orthogonal experts and designers can set objective parameters and weights through a GUI

The optimizer's objective function measures how well designs meet the designer's goals. For each new design task, designers may have different criteria, so this objective function can be customized in a simple GUI (Figure 14). In OPTIMUM, objective functions are constructed from a weighted set of objectives. Objectives are functions that take a design as an input, and output a value between 0 and 1 that estimates how well the design performs under some criteria. Designers weight objectives to express how important they are for a specific optimization. Given a set of objectives, $o \in O$, with weights β_o, the customized objective function will be the weighted sum of each objective's score (Equation 1a). Often, we will compare objective scores based on their percentage of a perfect score (Equation 1b).

$$f(d) = \sum_{o \in O} \beta_o o(d) \tag{1a}$$

$$F(d) = 100 \frac{f(d)}{\sum_{o \in O} \beta_o} \tag{1b}$$

Consider some cookie objectives. Just as butter effects the melting point of the batter, Chef Alton also explains that "baking soda reduces the acidity of the batter, thus raising the temperature at which the batter sets". Thus, Steve creates an objective that measures the target acidity level based on the proportion of baking soda to baking powder. Along with other objectives, this can be used to evaluate the texture of a cookie.

Objectives will be evaluated with each iteration of the optimizer and should be quick calculations or look-ups. An objective cannot, for example, run a lengthy simulation or request feedback from a user. As we will demonstrate, orthogonal experts can often define objectives for a variety of important criteria that designers can combine to support a variety of optimization tasks. However, this limits OPTIMUM in domains where objectives relate to aesthetics or cannot be measured efficiently.

8.2.2 *Modifiers*

Alongside an objective function, the optimizer needs strategies for improving designs. Modifiers are functions that take in a design and produce a new, slightly different, design. Modifiers provide small steps through the search space. Consider four simple cookie modifiers: two change the type of fat used (e.g., butter, shortening), and two that increase/decrease the baking soda. These and similar modifications to recipe parameters allow for full exploration of the cookie search space.

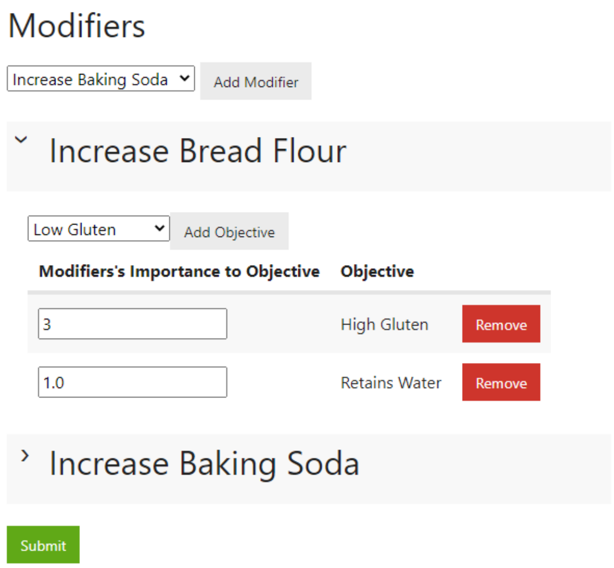

Figure 15: Non-programmer orthogonal experts and designers can assign heuristic weights in a simple GUI

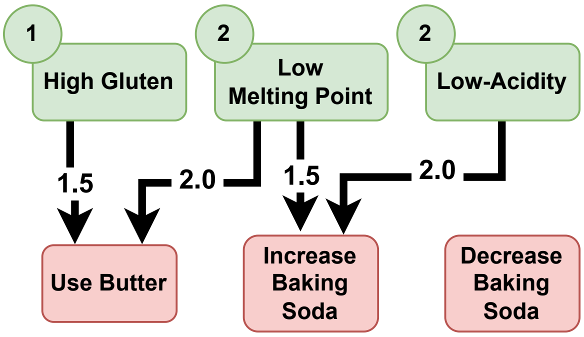

Figure 16: Example mapping of objectives to modifiers.

8.2.3 *Heuristic Maps*

Heuristics are strategies for improving a design with respect to a specific objective. In OPTIMUM, heuristics match a modifier to an objective that the modifier is likely to improve. Roughly, how likely the modifier is to improve that objective is expressed as an heuristic weight. That is, a modifier, m, that should improve an objective, o, has a heuristic weight $\alpha_{m->o}$. Building on prior work [106], we make the relationship between objectives and modifiers explicit in a *heuristic map* (Figure 16). Each modifier in a heuristic map can be mapped to multiple objectives and vice versa.

In another cookie example, Chef Alton expresses the value of switching to butter (modifier) for lowering the melting point (objective) with a heuristic weight of 2. Similarly, he relates increasing baking soda (modifier) to lowering the melting point (objective), but suggests that it is slightly less useful with a heuristic weight of 1.5. Note that the weights do not need to be finely tuned to express this logic.

$$E(d', m) = \sum_{o \in O} \alpha_{m->o} \beta_o (1 - o(d)) \tag{2}$$

Heuristic weights provide a domain agnostic way of measuring the expected value of applying a modifier to a design. The expected value is the sum of potential increases in objectives' scores between a design d and the resulting design from applying a modifiers, $m(d) = d'$, multiplied by the heuristic weight (Equation 2). This measure will be used by modifier selectors as a domain-agnostic measure of which modifier to apply in each optimizer iteration.

8.2.3.1 *Automatic Derivation of Heuristic Maps*

Algorithm 2 An algorithm that generates sample designs by iterative application of modifiers and uses them to set heuristic weights between objectives and modifiers.

Input seeds**:** Sample designs provided by orthogonal experts
Input $\mathbb{O}$**:** Objectives
Input $\mathbb{M}$**:** Modifiers
Input Rounds**:** the number of iterations to generate samples
Output $\mathbb{H}$**:** The estimated heuristic map

```
current_round ← seeds
designs_by_modifier ← {}                          ▷ Designs created by each modifier
o_increases_by_m ← {}      ▷ Counts of increases between o(d) and o(d′) given d′ = m(d).
for round ∈ Rounds do
    next_round ← {}
    for d ∈ current_round do
        for m ∈ 𝕄 do
            d′ ← m(d)
            add d′ to designs_by_modifier[m]
            add d′ to next_round
            for o ∈ 𝕆 do
                if o(d) < o(d′) then
                    increment o_increases_by_m[o][m]
    current_round ← next_round
ℍ ← {}
for m ∈ 𝕄 do
    for o ∈ 𝕆 do
        α_{m−>o} = P(o↑|m)                                 ▷ See Equation 3
        if α_{m−>o} > 0 then
            ℍ[o][m] ← α_{m−>o}
return ℍ
```

$$P(o_{\uparrow}|m) = \frac{|d' \in D_m | o(d) < o(d')|}{|D_m|} \tag{3}$$

OPTIMUM includes a tuning algorithm for learning heuristic weights from sample designs. We estimate heuristic weights given a set of designs, D_m, generated by applying a modifier to a previously generated design. We set the heuristic weight to the proportion of new designs that increased the objective score over all designs generated by the modifier (Equation 3). We generate this set of designs from seed designs. Over multiple rounds, we apply each modifier to the seed designs, and track increases in each objective's score. The generated designs seed the next iteration (Algorithm 2).

8.2.4 *Role of Orthogonal Experts and Programmers*

Building the heuristic library lays the foundation for domain-specific optimizers. An orthogonal expert and programmer must collaborate to build this library, which they can then reuse to create and experiment with different optimizers. The orthogonal expert directs the programmer, telling them what objectives are important, how to evaluate them, and what modifiers are useful in the domain. Orthogonal experts can then make the relationship between modifiers and objectives explicit by creating a heuristic map in a simple GUI. They can also generate a heuristic map by seeding the tuning method with their exemplar designs from the domain.

8.3 CONFIGURING METAHEURISTICS

Metaheuristic optimization covers a wide variety of optimization methods that can be applied to diverse domains without restrictions on their continuity or convexity. Further, they build on orthogonal experts' design strategies (i.e., heuristics). Most metaheuristic methods are characterized by how designs are chosen and modified in each iteration. Simulated annealing, for example, is defined by how designs are selected in each iteration, but design modification is left up to the programmer. Monte-Carlo methods, on the other hand, are defined by the probabilities of moving from one state to the next (i.e., modifier selection), but usually provide no stochastic method for jumping around the discovered design space. Making these different strategies pluggable enables programmers to mix and match methods until they find one suited to the domain.

The following section presents OPTIMUM's domain-agnostic library which defines a optimizer's meta-heuristic strategy. In OPTIMUM, metaheuristics consist of three components: *design selectors*, *modifier selectors*, and*stopping criteria*,. Stopping criteria determine when to stop iterating. A design selector chooses a design to modify in an iteration. A modifier selector determines which modifier to apply to that design. Each of these components make their decisions based on the objective function, heuristic map, and information collected during optimization in a *design population*. Our domain-agnostic library enables programmers to implement a wide variety of heuristic, stochastic, and metaheuristic methods, such as those used in our literature survey.

8.3.1 *Design Population*

Table 10: The types of information maintained by an optimizer's Design-Population. These values are used by stopping criteria and design selectors to determine the outcomes of each iteration of the optimizer.

Value	Key to Generated Designs
Iterations	Iteration which generated the design.
Scores	Objective function score of design. ($f(d)$, see Equation 1a)
Score Differences	Difference in score from prior design (i.e., $f(d') - f(d)$).
Objective Scores	Individual Objective scores of design. (i.e., $o(d) \forall o \in \mathbb{O}$).
Objective Score Differences	Differences in objective scores from prior design. (i.e., $o(d') - o(d)$)
Modifier Used	Modifier used to generate design.

As the optimizer generates new designs, the design population organizes information that helps determine when to stop and what designs to select in a domain-agnostic format. Designs are organized and sorted by: how well a design performed under the objective-function and individual objectives, the change from the prior design, the iteration in which the design was generated, and the modifier that generated it. Table 10 summarizes the information tracked by the design population.

8.3.2 *Stopping Criteria*

Table 11: Programmer parameterizable stopping criteria

Stopping Criteria	Programmer's Parameters	Definition
Exhausted Iterations	I: maximum iterations	Returns true when the I is reached.
Matched Threshold Score	N: Design count s: threshold score	Returns true when N designs have been generated with Objective Function scores $\geqslant s$
Objective Function Converged	I: iterations Δ: score changes	Returns true if I iterations have passed with no more than Δ difference in Objective Function Scores

We provide three parameterizable stopping criteria that determine when the optimizer should stop iterating and return the results to the designer (Table 11). We derive these from common stopping criteria across our survey of generative design and fabrication literature (see Chapter 2). Each stopping criteria is a function that, given the design population, will return true if the optimizer should stop. It is trivial to construct a variety of more complex stopping criteria by logically combining the results of these categories. For instance, we may want to continue searching for a local-maxima after a threshold score has been met. To do this, we can combine the results of a threshold and convergence stopping criteria so it only stops when both criteria are met. Further, we can set the optimizer to stop at a maximum number of iterations, returning the value of the combined stopping criteria (i.e., threshold and convergence) or an exhausted iterations stopping criteria. Stopping criteria can also be used to shift metaheuristic strategies mid-optimization, as they provide a mechanism for detecting relevant changes in the design population.

8.3.3 *Design Selectors*

Design selection is critical to the optimization process. An optimizer that chooses high-scoring designs quickly climbs towards a local maximum (i.e., hill climbing), but can miss distant, higher-scoring regions of the search space. Alternatively, choosing a design randomly will jump to a new regions, but will not consistently climb to a region's local maximum. Standard optimization algorithms balance between selecting quality designs and randomly searching by considering factors such as how long the optimizer has been running, the quality of previously discovered designs, and how often each design has been visited. OPTIMUM's library of design selectors provides a general, domain-agnostic strategy for managing these trade-offs. This approach can represent a wide variety of metaheuristic methods.

Design selectors are functions that, given the design population, return a design to be modified. Programmers configure a design selector with three components: the value to sort the design population by (see Table 10), the direction to iterate through that sorted designs, and a probability of selecting a design. Given these values, a design selector considers each design in sorted order, and has a random chance of selecting the design, given a programmer-set probability function. This continues until a design is selected or the last design is reached and selected.

8.3.3.1 *Sorting the Design Population:*

How the design population is sorted will have the most influence on which designs are selected. Most often, we choose a design with a bias towards those that are performing the best. This is true for a variety of standard metaheuristic methods. For example, evolutionary methods tend to select designs to mutate based on their quality [65]. This is done by sorting the designs from highest to lowest objective-function scores. Programmers can further refine this by searching based on specific objectives, which helps explore the Pareto boundary and can support methods such as Guided Local Search [256]. Other metaheuristic methods (e.g., Tabu search [81], ant-colony optimization [58]) select designs based on their visitation history. For these cases, the population can be sorted based on how many iterations selected the design. Similarly, we can bias the selector towards recent designs by sorting by the iteration that generated each design (e.g., iterative local search [158]).

8.3.3.2 *Selection Probability:*

Once the designs have been sorted, we need a selection probability. The simplest way to set this probability is a static value between 0 and 1. Some metaheuristic methods use more complex probability functions. For example, simulated annealing increases the probability of selecting high-performing designs as the number of iterations increases, causing it to converge on a high scoring region. The rate that the probability changes is often called a *cooling schedule*, following the metaphor that the optimizer is cooling like an annealing metal alloy. We provide a variety of parameterizable cooling schedules based on standard methods [4]. Note that these cooling methods are independent of how the population was sorted. We could, for example, cool based on the number of high-scoring designs, the top-score of the population, or convergence on a local-maxima. In most examples, we use an exponential multiplicative cooling function

[135] that increases from 0 towards 1 as the number of iterations increases; that is, the probability of selecting a design is $P(d) = 1 - S^i$, where S represents an entropy parameter between 0 and 1, and i is the current iteration count of the optimizer.

8.3.3.3 *Shifting Selection Strategies*

Programmers can shift design selectors mid-optimization with stopping criteria. The optimizer starts with one design selector and swaps selectors when a stopping criteria is met. For example, variable-neighborhood search [180] will perform a hill climbing process by modifying the last-generated design (design selector) until scores converge on a local-maxima (stopping criteria). Then, designs are selected randomly until some number of iterations is exhausted, triggering a shift back to the original design selector. Customizing design selectors using stopping criteria gives programmers even more flexibility in how designs will be selected at different stages of optimization.

8.3.4 *Modifier Selectors*

While our design selectors offer a variety of strategies for jumping around the discovered search space, the optimization process is only as diverse as the designs that are generated by applying modifiers to previously generated designs. Modifiers define the neighborhood of designs around a generated design. Because of this, choosing a modifier with each iteration is independent from design selection.

Modifier selectors are functions that return a modifier to generate the next design. Modifier selectors use the same pattern as design selectors. Modifiers from the heuristic map are sorted and considered by the selector given a probability function. Modifiers can be sorted by their known value, expected value, and application history. Just like design selectors, modifier selectors can be switched out given different stopping criteria.

8.3.4.1 *Known Value:*

Metaheuristic methods often step from a design to a high-scoring neighboring design. To sort by known-modifier value, we build a neighborhood set of designs by applying each modifier from the heuristic map. We then sort the modifiers based on how much each neighbor improved the design.

8.3.4.2 *Expected Value:*

When applying all modifiers would be too time consuming, we can sort modifiers based on their expected value (see Equation 2), estimated with heuristic weights. The expected values can serve as the probability of state changes within methods such as Monte-Carlo Markov chains.

8.3.4.3 *Modifier-History:*

The design population tracks both the changes in objective scores and which modifiers are used to create designs. From this information, we can keep track of how often, and by how much, each modifier has affected design scores. Just as we did in the heuristic map tuning algorithm (see Algorithm 2), we can use this information to estimate the probability that a modifier will improve the design (see Equation 3). This sorting criteria is useful for implementing methods like Ant-Colony optimization, where we want to bias modifier selection towards those that have a history of improving designs.

8.3.5 *Role of the Programmer*

Configuring the metaheuristic strategy of an optimizer is left up to the programmer, since it does not require orthogonal expertise. The programmer configures an optimizer by choosing a stopping criteria, design selector, and modifier selector. Programmers choose these components from OPTIMUM's library and can refine their behavior by adjusting the function's parameters. Figure 17 shows a simple GUI for

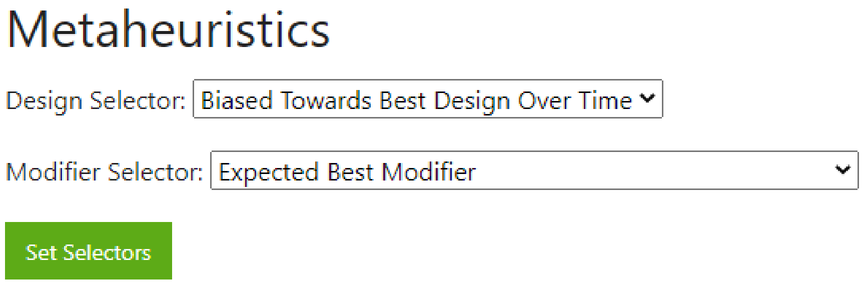

Figure 17: Simple GUI for adjusting metaheuristics.

switching out selectors, which gives curious orthogonal experts and designers a limited ability to adjust metaheuristics. However, they cannot adjust selector parameters without programming.

By default, OPTIMUM is configured to do simulated annealing, the most common metaheuristic method in our literature survey. The design selector sorts the population by objective function scores, and increases the probability of selecting high-quality designs over time. The modifier selector has an 85% bias for modifiers with higher expected values. The stopping criteria converges when 10 samples have stayed within a 5% difference of objective function scores. This optimizer utilizes a variety of the customizable properties of OPTIMUM's library for programmers to reference.

8.4 DESIGNING WITH OPTIMIZERS

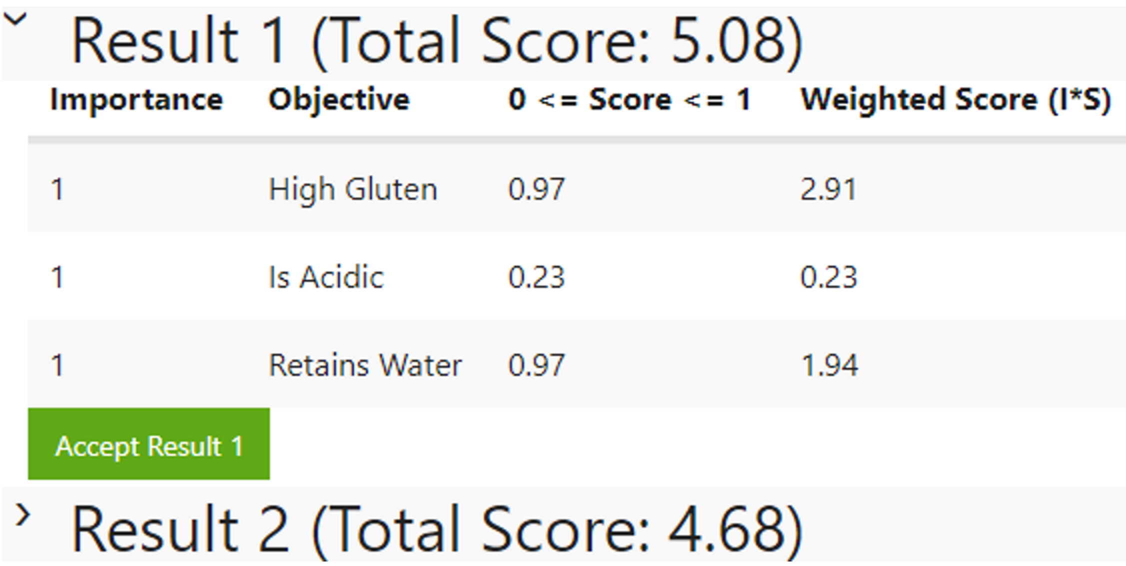

Figure 18: GUI for displaying optimization results.

OPTIMUM amplifies the programmer and orthogonal expert's expertise so that designers can independently use the optimizers. We assume that designers are also orthogonal experts; at a minimum, they can understand the objectives provided and how they can be combined to express different sets of design requirements. We do not assume that orthogonal experts have access to programmers, and thus use the optimizer entirely through the GUI. Just like a orthogonal expert, the designer uses the GUI to customize their objective function and heuristic map. The designer then starts the optimizer, which will generate a design population. After this, the GUI will display the resulting designs, sorted by objective function score. The designer can examine each design in the population (Figure 18). Being able to customize the objective function and heuristic map through a GUI gives designers significant control over the optimizer.

8.5 DEMONSTRATIONS

We have used OPTIMUM to build five domain-specific optimizers, two of which have been published in prior work [106, 108]. Each domain is built on the orthogonal expertise of non-technical experts, which we sourced by analysing design patterns from online communities Chapter 10, collaborating with researchers and users in the domain Chapter 9, and building on prior work with medical makers Chapter 3. In this section, we discuss three preliminary systems that demonstrate OPTIMUM's guiding principles. The cataract lens optimizer shows how OPTIMUM empowers an orthogonal expert (e.g., an Ophthalmologist) to participate in building an optimizer. The splint and tile-decor optimizers demonstrate OPTIMUM's flexibility.

8.5.1 *Cataract Lens Selection*

For our first demonstration, we recruited David, an Ophthalmologist, and Brian, a programmer, to build an optimizer that selects prosthetic lenses for cataract surgery [1]. Over two weeks, they coordinated by email to produce their optimizer. They demonstrate that programmers and orthogonal experts can collaboratively build optimizers with OPTIMUM that solve domain-specific problems.

As David explained, choosing a lens depends on a variety of factors, such as the patient's preferred outcomes, available lenses, and error tolerances. Cataract lens are defined by an a-constant, dependent on lens models, and an inter-ocular-lens power (IOL), which comes in 0.5 increments. Brian implemented this design representation as a simple Python Class with two parameters. David explained that he evaluates lenses with standard formulas [100, 209] that predict the patient's resulting refraction value (i.e., their glasses prescription). David weights these formulas based on patient demographics, which effect their accuracy. Brian encoded these formulas as objectives that compare the predicted refraction to David's target refraction. Brian implemented three modifiers that chose a random a-constant from available models, and incremented/decremented the IOL. After some trial and error, Brian chose the default design selector and stopping criteria, and chose a modifier selector that sorts by known value.

Brian and David had no prior experience implementing metaheuristic optimizers. Despite their lack of experience, their optimizer selected appropriate lenses. David provided de-identified patient data and lens selections for 10 prior patients. We compared his choices and those made by the optimizer. OPTIMUM selected the same a-constant with 100% accuracy and the correct IOL with 80% accuracy. In the two samples where the IOL differed from David's decision, it was by a single increment. David explained that this was a reasonable margin of error similar to the differences between different ophthalmologists. Further, David's preferred lens was always the second highest scoring option.

8.5.2 *Occupational Therapy Splints*

In this demonstration, we show how OPTIMUM's modifiable objective functions support automatic customization of splints, similar to those we created with Occupational Therapist (OT)s in Chapter 3. We represent splints with a set of parameters from a standardized splint pattern. We provide modifiers that increment and decrement each parameter by 1mm. We then derived objectives based on the OT's expertise.

Fit to Patient: The smaller the difference between a splint's parameters and corresponding patient measurements, the better it will fit. Fit is critical for ensuring splints properly restrict movement.

Restriction: Splints restrict movement of specific joints to support healing. OTs estimate the restriction by the width of the *wings* of the splint. Increasing width increases restriction.

Durability: Splints are more likely to break where wider wings and cooling holes introduce material strain. OTs estimate durability with cooling density and inversely to the wing widths.

Comfort: Restriction and durability reduce comfort. While some discomfort is expected, a splint that is too uncomfortable is likely to be abandoned. As such, OTs consider the value of introducing cooling holes and narrowing the lower wing of the splint, which can cause discomfort when moving the wrist.

1 Participant selected pseudonyms

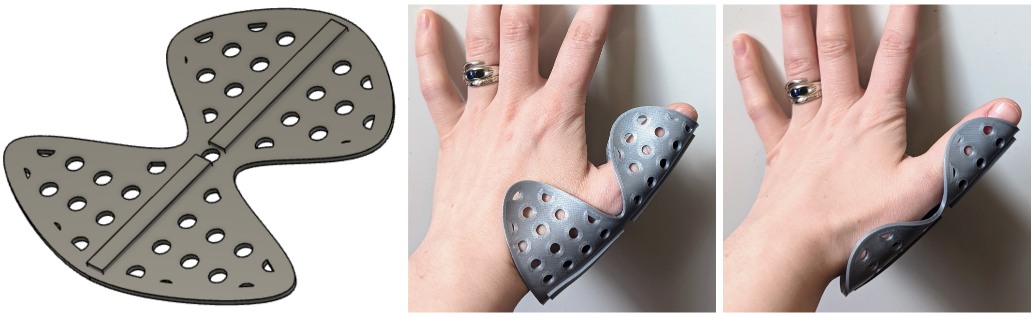

Figure 19: Splints are thermoformed to the patient's hand.

By changing the objective weights, we can quickly customize splints. Figure 19 shows two splints. The left splint has an increased weight on the restriction objective, which widened the wings. The right splint decreased this weight and increased the weight on comfort, causing the lower wing to narrow. Using this optimizer, OTs could easily experiment with different weights and produce a variety of splints to test out with patients, prioritizing different aspects of the design. This shows the value of OPTIMUM's re-configurable objective functions for end-user customization.

Heuristic weights play a significant role in optimizer results, because they express the relationship between objectives and modifiers. To demonstrate this, we compared three splint heuristic maps across 100 optimization trials, each with randomly weighted objectives. One heuristic map used weights derived from our study notes. The second used weights tuned by seeding our tuning method with the sample splints created by the OTs. The third used equal weights between all objectives and modifiers. Figure 20 shows how each of these heuristic maps performed. A Welch's Anova test shows a significant ($F = 439.79, p < 0.01$) effect on the splints' objective scores. A Games Howell Post-hoc analysis shows a significant difference ($p < 0.01$) between the expert and uniform ($T = 29.3$) conditions, as well as the tuned and uniform ($T = 29.2$), but not the expert and tuned conditions ($T = 0.9$). This shows the value of expressing domain-expertise directly through heuristic maps or indirectly through our tuning algorithm.

The splinting domain demonstrates the value of including orthogonal experts in heuristic strategies. While there may be a more efficient or effective optimization method, we were able to build simple and effective splinting tool using a description of the traditional splint design process. By structuring the optimization around heuristics, we can embed the OT's design process into a tool. This reduces a manual design process, that takes many iterations over many clinic appointments, to an optimization process that consistently converges in under 500 milliseconds. The resulting design is ready to print and provide to the patient. Further, by structuring the objective function around a set of weighted objectives, the OTs can quickly explore different design trade offs by adjusting the objective weights (e.g., durability vs comfort). Using a more complex and specialized optimization technique may not afford this flexibility.

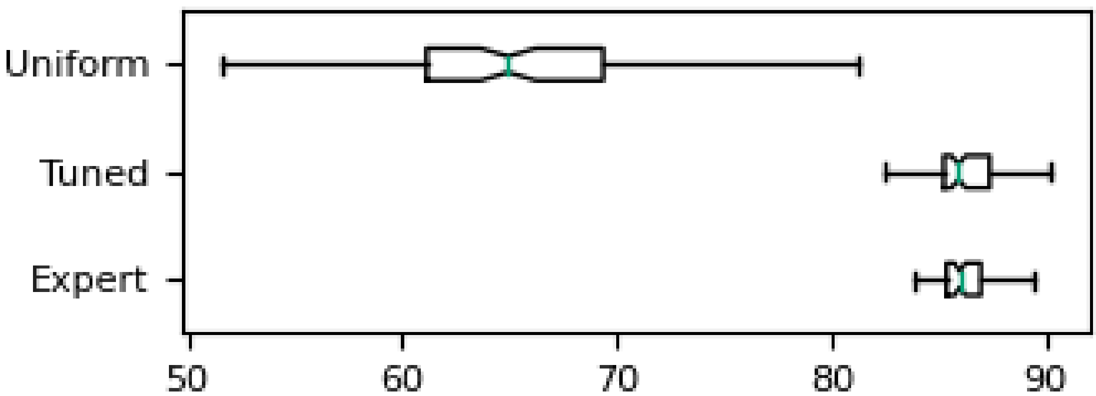

Figure 20: Percentage of perfect score across experiments.

Figure 21: Example tile-decor

8.5.3 *Replicating Tile-Decors*

This demonstration replicates Chen *et al's* prior work [44] that generates "*surface-like objects composed of connected tiles*". Their heuristic method has a structure similar to our optimizers. Their objective function sums four weighted objectives that minimize neighborhood distance, maximize surface approximation, maximize hinge-placement, and minimize repulsion. Each objective is associated with a modification that either randomly places a tile, attracts tiles together, repulses tiles, or scales tiles to their neighbors. We have replicated this method without OPTIMUM, and then created seven OPTIMUM optimizers that use different metaheuristic strategies without changing the underlying heuristics. We examine the differences between metaheuristics on resulting scores and convergence times.

Plugging in different design and modifier selectors tests metaheuristic strategies, without modifying the domain-specific code used to evaluate and generate tile decors. We conducted seven experiments (Table 12) where we tasked the optimizer with generating a packed, cylindrical tiled surface (Figure 21). For each optimizer, we ran 100 optimizations with random weights on the different objectives. We collected the resulting objective score and the time it took to converge. We kept the stopping condition constant, halting the optimization when a tile-decor was discovered that achieved 95% of a perfect objective score or when 10000 iterations was exceeded.

Table 12: Optimizer configurations. A-D vary design selector. D-G vary modifier selector. G was the best optimizer

Optimizer	Design-Selector	Modifier-Selector
A	Select random design	Select modifier with highest expected value
B	Select highest scoring design	Select modifier with highest expected value
C	Select highest scoring design with 0.85 probability	Select modifier with highest expected value
D	Select highest scoring design with probability increasing with iterations	Select modifier with highest expected value
E	Select highest scoring design with probability increasing with iterations	Select best-known modifier for objective function
F	Select highest scoring design with probability increasing with iterations	Select best-known modifier for lowest scoring objective
G	**Select highest scoring design with increasing probability**	**Select best-known modifier for most important objective**

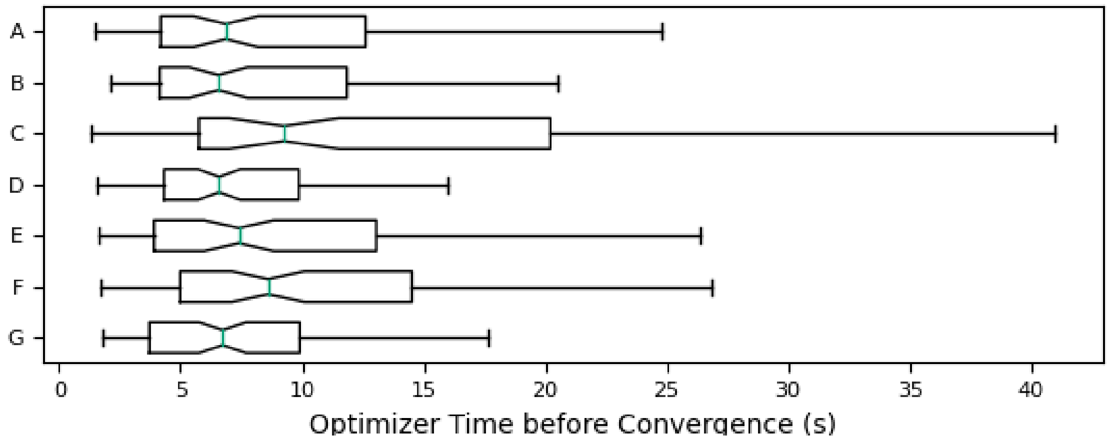

Figure 22: Comparison of tile-decor convergence times in seconds. See experiment descriptions in Table 12

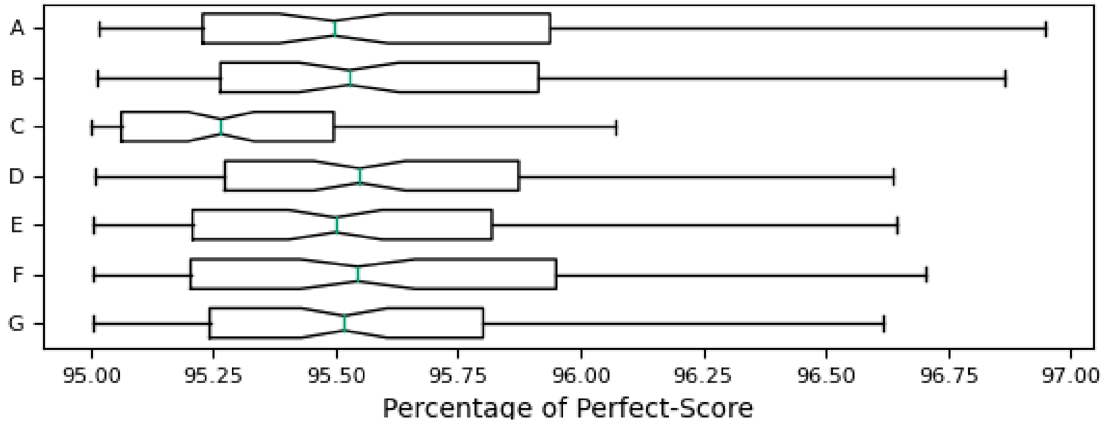

Figure 23: Comparison of tile-decor percentages of perfect scores. See experiment descriptions in Table 12

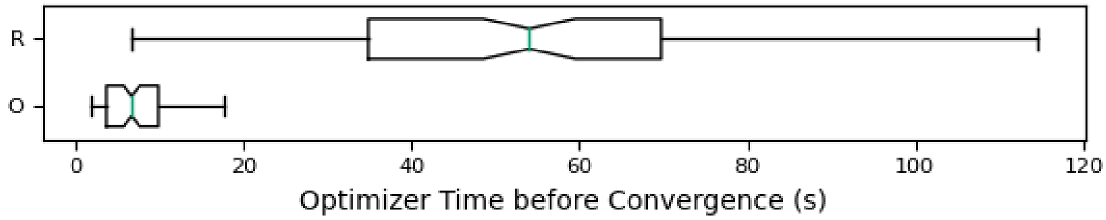

Figure 24: Comparison of convergence times between OPTIMUM (O) and replicated (R) optimizers.

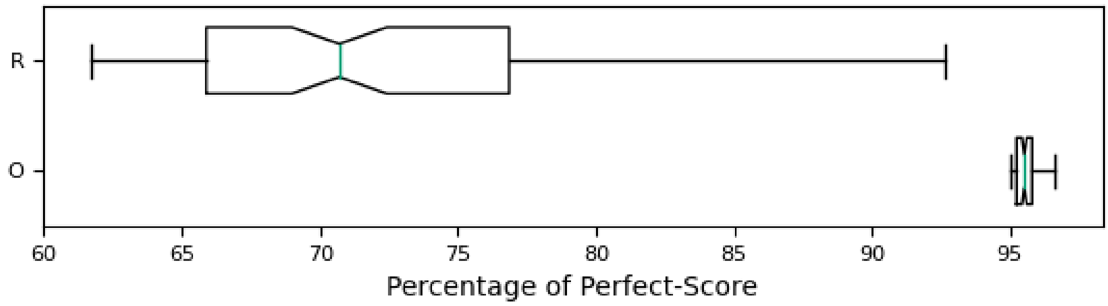

Figure 25: Comparison of scores between OPTIMUM (O) and replicated (R) optimizers.

Table 13: Results of Games Howel Post-hoc analysis comparing tile decor experiments. *Indicates significance (p<0.05).

	A	B	C	D	E	F	G
A Score		-0.2	*4.8**	-0.1	0.6	0.2	0.6
A Time		0.6	*-3.02**	0.8	-0.4	-1.4	1.5
B Score	-0.2		*5.3**	0.1	0.8	0.4	0.8
B Time	0.6		*-3.5**	0.2	-1.0	-2.0	1.0
C Score	*4.8**	*5.3**		*-5.6**	*-4.4**	*-4.8**	*-4.9**
C Time	*-3.02**	*-3.5**		*3.6**	2.6	1.7	*4.1**
D Score	-0.1	0.1	*-5.6**		0.8	0.4	0.7
D Time	-0.4	-1.0	2.6		-1.1	-2.1	0.7
E Score	0.6	0.8	*-4.4**	0.8		-0.4	-0.1
E Time	-1.4	-2.0	1.7	-2.1		-0.9	1.8
F Score	0.2	0.4	*-4.8**	0.4	-0.4		0.4
F Time	-1.4	-2.0	1.7	-2.1	-0.9		2.7
G Score	0.6	0.8	*-4.9**	0.7	-0.1	0.4	
G Time	1.5	1.0	*4.1**	0.7	1.8	-1.4	

The first four experiments vary the modifier selector (Table 12 A-D). A Welch's Anova showed the significant effect of varying the modifier selector ($p < 0.01$) on scores ($F = 17.4$) and optimization time ($F = 4.6$). We then varied design selectors (Table 12 D-G), and a Welch's Anova showed no significant effect on score ($p > 0.1, F = 0.3$), but did find an significant effect on time ($p < 0.05, F = 3.1$). A Games Howell post-hoc analysis shows the differences across experiments (Table 13). We selected G as our final optimizer because it converged fastest (Figure 22, 23).

Testing different combinations requires no new code, only a change to the optimizer's configuration. In comparison to our replication of Chen *et al's* original algorithm (Figure 24, 25), the OPTIMUM optimizer converged on significantly ($p < 0.01$) higher scores ($T = 39.8$) and converged significantly faster ($T = -16.12$), based on a Welch's T-test. Further, experimenting with different metaheuristics gave us unique insights into the optimization problem. Unlik, Chen *et al's* staged optimization method, our optimizer switched between increasingly different objectives with each iteration. This revealed a tension between attraction and repulsion objectives that is poorly supported by optimizing them in different stages.

OPTIMUM gives programmers flexibility when implementing and selecting metaheuristic strategies, while still reusing their domain-specific code. This supports rapid prototyping of optimizers. The flexibility of the system makes it simple to configure satisfactory optimizers and to compare a variety of optimizers.

8.6 DISCUSSION

OPTIMUM offers a new way of implementing generative design systems. Optimization requires orthogonal expertise, and OPTIMUM gives orthogonal experts a format to make their expertise explicit. Orthogonal experts can recognize what features of a design need to be evaluated, ways of calculating those evaluations, and ways to modify designs to meet those goals. In many cases, orthogonal experts may intuitively know the relationship between some of these objectives and modifiers and can express them in a heuristic map. When these relationships are too complex or cannot be intuited, we can automatically learn them from the expert's sample designs. OPTIMUM provides a flexible framework for defining domain-agnostic metaheuristics that operate on domain-agnostic representations of heuristics (i.e., heuristic maps) and the state of an optimization (i.e., design population). This flexibility helps programmers to rapidly prototype optimizers in the domain. In effect, the optimizer amplifies the optimization expertise of the programmer by making the optimization accessible without programming, as well as the orthogonal expert's expertise by applying their design heuristics to a wide variety of domain-specific problems.

In this chapter, the demonstrations describe two simple tools that build on the expertise of clinicians (e.g., David the Ophthalmologist and the OTs). These are examples of clinical CAD tools that can be produced with OPTIMUM. Both demonstrate the value of small adaptations to standardized design and amplify the expertise of the clinicians we worked with. Further, they show the clinicians' emphasis on production, creating designs that are ready for patients to use, rather than prototyping novel medical devices. Tools like these are complementary to generalized prototyping tools, like PARTs (Chapter 7), and traditional CAD tools.

8.7 LIMITATIONS AND FUTURE WORK

OPTIMUM has limitations that open opportunities for future work. First, while we have reduced requirements for programmers' orthogonal expertise, programmers still need to collaborate with orthogonal experts to craft design representations, objectives, and modifiers. OPTIMUM's GUI enables non-programmers (i.e., orthogonal experts, designers) to modify objective and heuristic weights, however, there remain opportunities to define design-representations, objectives, and modifiers without programming. While this may not be feasible across all domains, interfaces for visually programming parameterized designs or graph structures could build on OPTIMUM's core optimization library.

8.8 CONCLUSION

OPTIMUM supports a new approach to implementing generative design tools in unique domains, built around guiding principles for clinical CAD tools. Unlike other optimization toolkits, OPTIMUM separates the roles of programmers who implement optimizers and orthogonal experts who guide optimizers to good results. We do this by having orthogonal experts and programmers collaboratively build a library of domain-specific objectives and modifiers, which orthogonal experts can associate to define heuristics.

The next two chapters will describe in greater detail two additional generative design tools implemented with OPTIMUM. The first, Maptimizer, helps people who are blind or have low-vision (BLV) generate tactile maps that are tailored to their unique needs and abilities. Unlike the demonstrations with Ophthalmologists and OTs, Maptimizer positions BLV people as orthogonal experts in their own navigational practices. This demonstrates that the principles behind clinical CAD tools extend from clinical to assistive applications. Chapter 10 uses these principles in a new domain, automatic machine knitting. While not strictly a medical making domain, knitting is a space that benefits from tools that enable pattern reuse, support small adaptations to garment designs, focus on manufacturing and production, and amplify the expertise of knit-designers and hand knitters. These demonstrations show some of the breadth of these principles through novel, domain-specific generative design tools.

9 MAPTIMIZER: TAILORING TACTILE MAPS TO USERS NEEDS

Tactile Maps display information about a geographic location using a set of raised tactile features that can support navigation by people who are blind or have low-vision (BLV). When used alongside other navigation aids, a tactile map can help the user better understand a geographic space and the relationship between different geographic features. For instance, a tactile map could help a user identify a geographic feature that can support them in orienting themselves, such as a loud fountain at the center of a University courtyard. However, to better support navigation, tactile maps must be more widely available and include information that is tailored to the user's specific needs, their abilities, and the mapped location.

The availability of tactile maps is limited by the inaccessibility of tools for generating tactile maps [86] and lack of accessible customization for users [246]. The increasing availability of consumer 3D printers offers an opportunity to produce complex tactile maps using a variety of approaches (e.g., [86, 110, 111, 246]). In a few cases, BLV users are afforded some customizing opportunities, ranging from setting the location and scale of a map [126], to individually tailoring how different geographic features are represented [246].

The dichotomy between non-customizable interfaces and overly customizable but burdensome ones leaves users struggling to create maps that meet their needs. As in many assistive domains, customization is critical, however, it also creates a complex design task. Ideally, the user would be able to rapidly iterate and create a variety of maps to meet their needs. However, such iterations are largely inaccessible because 3D printing maps is resource intensive, and many users may not consider this extensive design task worth their efforts. Further, the user alone may lack critical contextual information about the unfamiliar location being mapped. Tactile map generation requires rapid iteration over a variety of designs that incorporate the user's context, as well as geographic context.

Maptimizer is a generative design tool implemented with OPTIMUM (Chapter 8) that gathers user's preferences from a screen-readable web application, and generates a tactile map that optimizes the communicability of the tactile representations and the informativeness of the geographic features embedded in the map while minimizing map clutter. When creating a tactile map, Maptimizer adjusts how geographic features are represented based on what the user prefers and can best distinguish given their abilities. Further, Maptimizer optimizes maps based on contextual information about the location.

9.1 BACKGROUND: TACTILE MAPS

Previous work demonstrates the potential benefits and challenges of using tactile maps as a way-finding and orientating tool for BLV users [67, 123]. While access to rich information raises the comfort level associated with independent travel and orienting, it is important to determine which presented information will be most helpful for the user to avoid unnecessary information overload [92]. Williams *et al* note the distinct needs of BLV users, noting that they vary significantly based on a user's preferences (e.g., traveling with a cane or guide dog) and the context (e.g., outdoors, indoors, crowded, open) [266]. These unique differences inform the planning of routes [129] and travel aids [199].

9.1.1 *Tactile Map Generation*

Several researchers have tackled the challenges of making tactile maps more useful, easier to fabricate, or inexpensive, using a multitude of algorithms and models for selecting what information to include and how to portray it [258]. Wang *et al* created a system that takes a map (e.g., from an online site) and makes an SVG rendering of it for 3D printing with Braille labels [259, 260]. Taylor *et al* introduced the idea of customization by generating 3D printable maps with a small set of user-customizable options [246]. They note the challenge of enabling advanced customizations without making the interface too complex or inaccessible to screen readers. Miele *et al* found that braille labels, which can quickly clutter the map, can be replaced with embossed overlays that played audio labels as features when touched [178]. Gotzelmann *et al* expanded on the use of this capacitive sensing technique by 3D printing the maps [86]. Giraud *et al* found that memory retention of information about points of interest was improved when text was shared on-demand from a user touching a tactile map feature [78]. Similarly, Taylor *et al* created tactile maps which would be placed over the phone in a 3D printed case, both of which had buttons that afford interaction with the map [248]. While Taylor *et al* [246] and Gotzelmann *et al* [86] produce 2.5D maps, Holloway *et al* designed maps with more complex 3D icons that mirrored the real life objects they represented [110, 111]. These icons were easily recognizable, and reportedly helped users develop mental models of the location. Each of these tools focuses on new ways of producing tactile maps, however, little focus is centered on the task of designing a customized map.

9.1.2 *Optimizing Map User Interfaces*

Metaheuristic optimization techniques are often used to render digital maps. For example, Agrawala and Stolte [10] use a simulated annealing optimization method to simplify how routing information is displayed to a driver so that it is easy to understand, as well as minimizing how much attention a driver needs to pay to a map while driving. Grabler *et al* [88] also applied simulated annealing to design tourist's city maps that present the most valuable information to individual users, rather than presenting the same information to all users. This demonstrates the use of these optimization techniques to create customized maps that meet the needs of individual users. Similarly, Lee *et al* [149, 150] use optimization techniques to select what features are displayed to the user, prioritizing features that provide the most important contextual information during route navigation. Individually, these systems present optimization techniques that focus on three key characteristics of maps: how communicative the representations of information are [88], how informative information displayed is [149, 150], and how much of the user's attention is drawn from their primary task (e.g., driving) to the map itself [43]. These characteristics are also critical to the design of high-quality customized tactile maps.

9.2 MAPTIMIZER

Maptimizer enables people who are BLV to independently generate a tactile map that is uniquely tailored to their preferences. In particular, we consider two types of preferences that will affect the quality and usability of the tactile maps. First, a user's preferences may affect what types of *geographic features* are presented in the map. Beyond simply showing roadways and buildings, a user may need specific information about points of interest, accessible infrastructure (e.g, stairs, ramps, tactile markers, audible and tactile traffic signals), general infrastructure (e.g., walkways, entrances to buildings), and/or amenities (e.g., benches, restrooms). Second, users may have preferences for how this information is represented (i.e., its *representation*), such as the types of symbols used to represent information and the size of those symbols. However, designing a tactile map requires more than a user's preferences. The information they prefer must be supported by sufficient contextual information (e.g., features critical for orientation) to navigate effectively, and showing too much information can clutter the map, making it difficult to read. Maptimizer is a generative design tool that uses optimization methods to generate maps that consider these different requirements and prioritize the end-user's preferences. We use a two stage optimization process, that first uses linear-programming to pair communicative representations to informative geographic features based on users' preferences, and then uses an OPTIMUM optimizer that tunes sizing parameters of representations to increase communicability and reduce clutter.

In Maptimizer, tactile maps are generated by pairing geographic features with representations of those features. To avoid user confusion, each representation can only be used for one geographic feature, and each geographic feature can only be shown with one representation. Representations will generate a component on the tactile map. For instance, a *path* representation will create a series of raised lines and a *location* representation will create a peg that sticks out of the map to mark that location (Figure 26). Similar to [87], these representations are used to present three types of *geographic features*: regions (e.g., buildings, water ways), paths (e.g., roads, sidewalks, bus routes), or locations (e.g., points of interests, benches, tactile

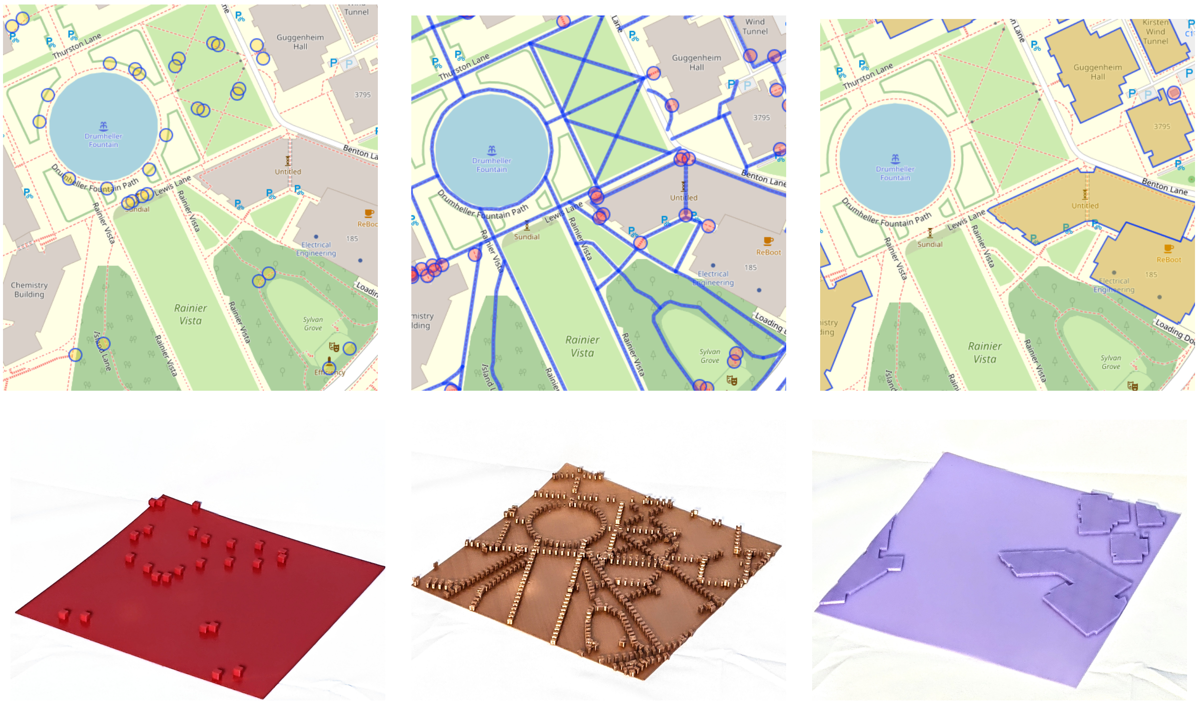

(a) Locations are represented with differently shaped pegs.

(b) Paths are a series of of raised dashed lines.

(c) Building Regions represented by raised polygons on the tactile

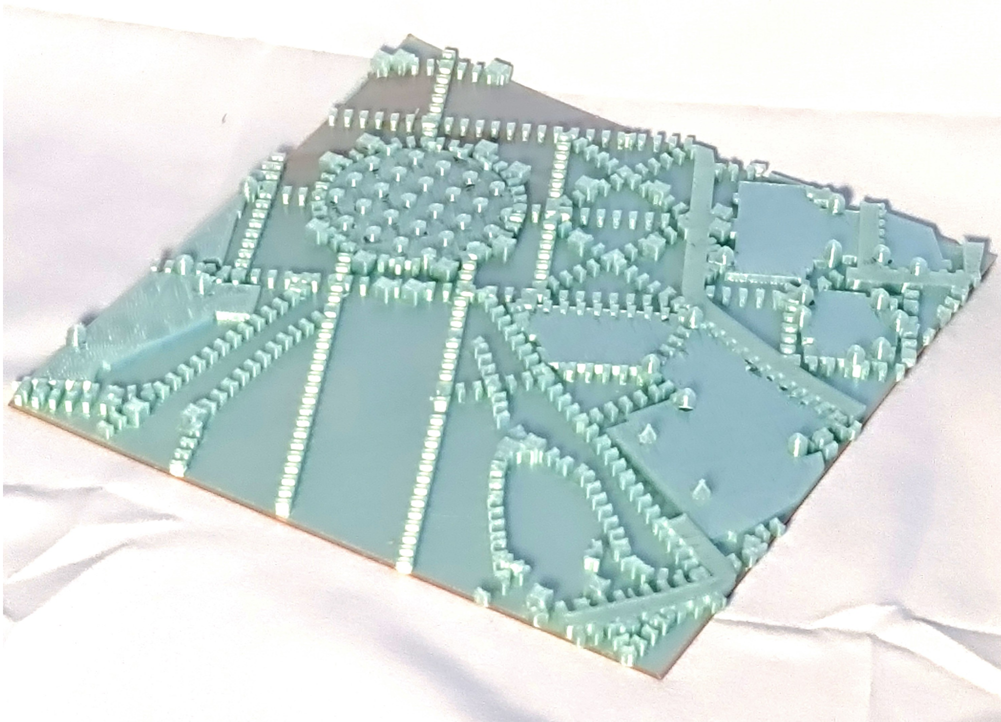

(d) Combinations of different representations are added onto a base-map

Figure 26: A tactile map pairs geographic features to tactile representations of those features. There are three types of features which can be related to specific representations: (a) locations, (b) paths, (c) regions. Combining different representations creates a full tactile map (d).

pavement). For example, a raised path could represent roads or footpaths but cannot represent a set of benches or buildings.

9.2.1 *Geographic Features*

We gather information about geographic features from the OpenStreetMap API [193]. Which geographic features are available depends on the location and data available through OpenStreetMap. Most areas have data about major architectural features (e.g., roads, buildings) and ecological features (e.g., green spaces, water ways). Some areas may also include accessibility features (e.g. accessible traffic signals, curb-cuts, tactile pavement). Our goal is to collect a wide variety of information from OpenStreetMap, which enables users to determine which pieces of information are most critical to them. They may consider points of interests, routes they intend to take, or accessibility features that will help them plan their trip.

There are two important attributes of a geographic feature: the user's preference for that feature and the area of the map taken up by the feature. We denote a user's preference for a geographic feature, g, as $u(g)$. This value is provided through a user interface and falls on a scale from 0 (not important) to 10 (most important). We denote the area of a geographic feature as Area(g). The area of a region-feature is the area of polygons that make up the region (e.g., the area covered by buildings). Path-features are connected lines, so we approximate the area as the length of these lines multiplied by a width of 1mm. Similarly, location features are a set of coordinates in the map, so we approximate their area as the count of locations. This essentially treats each location as a dot on the map that has an area of 1mm^2. Note that we do not consider sizing-parameters, such as the size of pegs that mark each location or the width of roads. These are accounted for by the size of representations.

$$||g|| = \sum_{f \in g} ||f|| \tag{4a}$$

$$||f|| = \begin{cases} \mathtt{area}(f) & \text{if } f \text{ is an Region} \\ \mathtt{length}(f) & \text{if } f \text{ is a Path} \\ 1 & \text{if } f \text{ is a Location} \end{cases} \tag{4b}$$

9.2.2 *Representations*

Regions, paths, and locations can each be represented in a variety of ways. Locations are represented by differently shaped (i.e., pyramids, domes, and cubes) pegs that stick out of the map at the corresponding coordinates. Paths are represented by raised solid, dashed (Figure 26.b), or dotted lines. For regions, the simplest representation is to emboss a set of polygons shaped like that region (Figure 26.c). To differentiate regions, textures can be added on top of these embossed polygons. In total, we include ten representations: four regions, three paths, and three locations. Our set of representations is extensible, and this set ensured a wide variety of map designs for our user evaluation.

These representations are individually parameterized. Each representation has a depth parameter, d_r, which determines how much it is raised out of the map base. Path and location representations, additionally, have a width parameter, w_r, which determines the width of the path or peg. These parameters will determine how much space is taken up by a represented geographic feature in the final tactile map. This, in turn, effects how easy different geographic features are to identify and distinguish. Larger representations will stand out of the map more. Correspondingly, smaller representations will be more difficult to find and identify. Additionally, users provide their own preference-ranking, $u(r)$, for each type of representation in the user interface on a scale from 0 to 10.

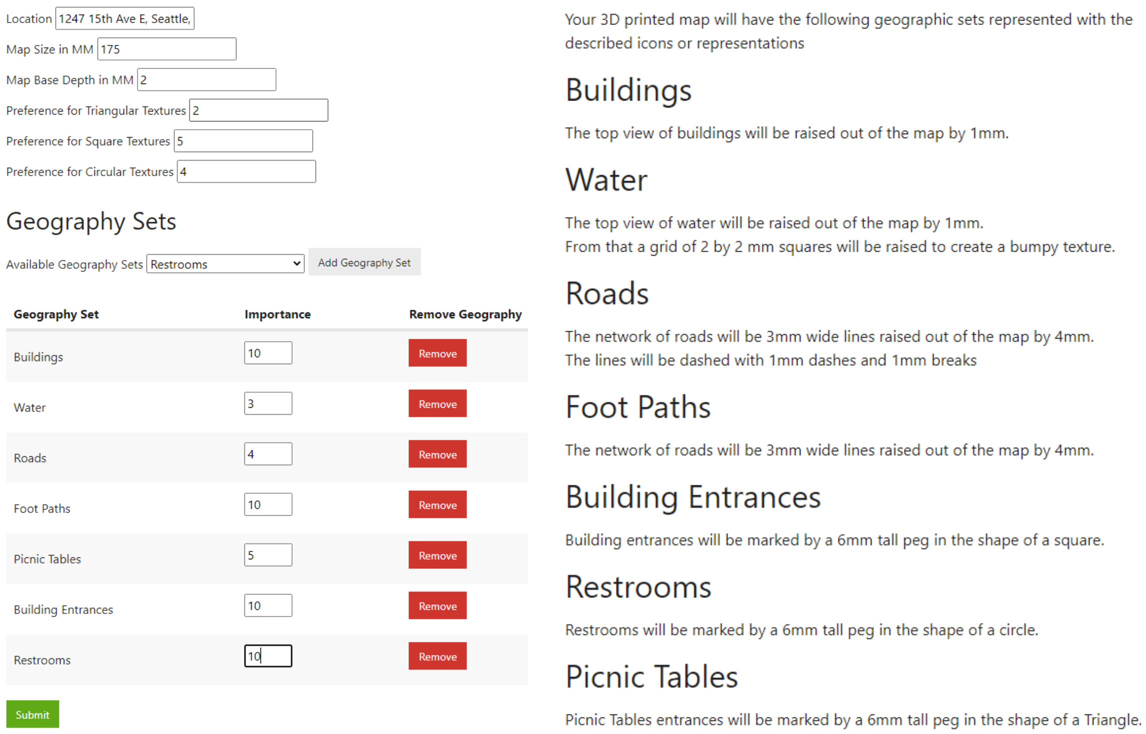

(a) Maptimizer Screen Readable Interface

(b) Screen Readable Optimized tactile map Legend

Figure 27: Users set their preferences for different types of geographic features and how representations are made in a simple screen readable interface. Maptimizer displays the optimized tactile map's legend and creates an STL file for 3D printing.

9.2.3 *User Interface*

We provide a simple, screen-readable web interface implemented with OPTIMUM that collects user's preferences for different ways of representing information and different geographic features (Figure 27). Once the user has provided the location of their map, either an address or coordinates, it queries the OpenStreetMap data set for geographic features in that area. Each of these geographic features is presented in a table (Figure 27a), and the user can rank them on a scale from 0 to 10, with 10 being the most important geographic features. These rankings are used in the optimization process and denoted $u(g)$. To generate each tactile map, the user also ranks different types of representations (i.e., $u(r)$) based on their textures (e.g., square, circular, triangular). Once the tactile map is generated, a screen readable legend describing the map is presented to the user (Figure 27b). Each included geographic feature is labeled with a header and is followed by a brief description of its representation. Maptimizer also outputs a 3D printable model of the optimized tactile map.

9.2.4 *Objective Function*

Tactile maps have three conflicting qualities: *communicability*, how well the map conveys information; *informativeness*, how valuable is the information the map conveys; and *attention-costs*, how cluttered the map is (a noted issue with tactile maps [87]). Lee *et al* [149, 150] define ways of measuring the communicability of each representation and use more-communicative representations with the most informative map features. Grabler *et al* [88] argue that the informativeness of a map is dependent on the geographic features it represents and how those align with the user's goals. Agrawala and Stolte [10] minimize attentions costs in the context of driving tasks, but this is also relevant in the low-resolution design space of tactile maps, especially if a user chooses to use the map while navigating the space. We have adapted these concepts to the domain of tactile maps.

Considered alone, each of these attributes is easy to optimize but will produce a poor performing map. For example, using only preferred representations will be very communicative, but may not show every piece of important information without cluttering the map. Then again, a sparse map with few key pieces of information will be less cluttered but will leave out important information. We formalize these properties (i.e., communicability, informativeness, attention-cost) to create the objective function which Maptimizer will maximize by generating unique tactile maps based on a user's preferences and a location's context.

For the purposes of optimization, we define a tactile map as a set of a representations paired to geographic features, such that each geographic feature has one, and only one, unique representation. We denote the pairing between a geographic feature, g, and a representation, r, $p_{r \to g}$, and the set pairs in a specific tactile map as $\mathbb{P}$. Representations and geographic features have user rankings (i.e., $u(r), u(g)$). $C(r)$ denotes a representation's communicability. $I(g)$ denotes a feature's informativeness. $A(r, g)$ denotes the attention cost of representing a geographic feature, by measuring how much clutter the pairing of r to g creates. Generally, for each pairing of representations and geographic features, we want to maximize the sum of the representation's communicability and the geographic feature's informativeness while minimizing their attention cost ($o(r, g)$, Equation 5a). We use the control variables ζ, ι, and α to weight communicability, informativeness, and attention costs in the objective function. For all maps produced for our user study, we set their values to $\zeta = 1$, $\iota = 2$, and $\alpha = 1$, based on initial pilot tests. To evaluate a whole tactile map, we calculate the sum, $O(\mathbb{P})$ (Equation 5b), for each pairing of representations and geographic features.

$$o(r, g) = \zeta C(r) + \iota I(g) - \alpha A(r, g) \tag{5a}$$

$$\text{Maximize: } O(\mathbb{P}) = \sum_{p_{r,g} \in \mathbb{P}} o(r, g) \tag{5b}$$

Internally, this objective function is composed of OPTIMUM objectives for communicability, informativeness, and attention-cost for each type of representation, geographic feature, and possible pairing. The control variables are objective-weights provided by the user. The objectives are implemented as functions that calculate the following values.

COMMUNICABILITY

$$C(r) = \begin{cases} u(r)\frac{d_r}{d_{max}} & \text{if } r \text{ represents a region} \\ u(r)\frac{d_r w_r}{d_{max} w_{max}} & \text{otherwise} \end{cases} \tag{6}$$

The communicability of a representation measures how easy a representation is to distinguish from other representations in the map. Our estimation of communicability considers the user's ranking of that representation, $u(r)$, and how large they are (i.e., larger representations are easier to distinguish). We estimate size as the proportion of the size of the representation, defined by its parameters (d_w, w_r), over the maximum allowed size of those parameters (d_{max}, w_{max} (Equation 6). For each representation, these parameters can range from 1mm to 10mm. Generally, larger representations with greater parameters, will be easier to recognize, making them more communicative. We multiply this by the user's ranking. A strong preference for a particular representation (e.g., high $u(r)$) increases communicability.

INFORMATIVENESS

$$I(g) = \frac{\text{Area}(g)}{\text{Map Area}} u(g) \tag{7}$$

The informativeness of a geographic feature is highly dependent on the user's information preferences, which, in turn, is dependent on how they expect to use the map. Additionally, informativeness is dependent on information which the user may not be aware of. For example, a user may highly value common features such as buildings and roads but be unaware of a water feature (e.g., a lake) which is critical to understanding a location. A highly informative map includes features that provides context, that a user prefers, and that prioritizes features that meet both criteria.

We estimate informativeness as the product of two terms (Equation 7). The first term approximates how critical a feature is to understanding a location based on how much of the map it takes up. We use a heuristic which assumes that if a geographic feature takes up a large portion of the map, it must provide significant contextual information. We calculate this as the proportion of the area of the geographic feature over the total area of the map. The second term is the ranking of each geographic feature, $u(g)$.

ATTENTION-COST

$$A(r, g) = \begin{cases} \frac{d_r \text{ Area}(g)}{\text{Map Volume}} & \text{if } r \text{ represents a region} \\ \frac{d_r w_r \text{ Area}(g)}{\text{Map Volume}} & \text{otherwise} \end{cases} \tag{8}$$

Alone, maximizing communicability and informativeness would produce tactile maps packed with many communicative representations of informative geographic features. However, these tactile maps would be dense and cluttered, and may confuse users or obscure valuable information. To penalize cluttering the tactile map, we subtract an attention-cost term, which estimates how much of the tactile map is covered by the representation of a specific geographic feature. To measure this, we calculate the *volume* of a representation-geography pair by multiplying the sizing parameters of the representation (e.g., d_r, w_r) by the area of the geographic feature. We divide that by the *map volume*, which is simply the map area multiplied by the maximum allowed depth of any representation. The proportion of the represented geographic feature's volume to the total map volume will penalize larger representations and geographic features that might occlude the rest of the map.

9.2.5 *Multi-Stage Tactile Map Optimization*

Maptimizer generates tactile maps in two optimization stages (Figure 28). The first stage determines what types of representations will be paired with geographic features. It does this by considering the user's preferences (e.g., $u(r), u(g)$) gathered from the Maptimizer interface. The second stage optimizes the sizing parameters of each of the representations that were selected in the first stage. Tuning these parameters makes the geographic features easier to identify (i.e., communicative), while minimizing clutter (i.e., attention-cost). Our optimization method is not guaranteed to find the globally optimal tactile map. However, based on the results of our user evaluation, we expect that the space of tactile maps has many high-quality local-maxima.

9.2.5.1 *Stage 1: Representation and Geographic Feature Pairing*

The first step to generating a tactile map is to identify the most informative geographic features and pair them to representations that will best communicate that information. The optimization starts with a pairing of geographic features to representations that would maximize our objective function (Equation 5b) if there were no constraints. This would be a set of pairs where every geographic feature is paired once to every representation, meaning that all of the most informative features are paired with all of the most communicative representations. However, this map would be unusable because there would be no way to distinguish between different features on the map. Thus, we apply three constraints. First, each representation can only be paired with one geographic feature. Second, each geographic feature can only be paired to one representation. Finally, a representation can only be paired to a geographic feature if they are *compatible*. Naturally, region representations are compatible with geographic regions, path representations are compatible with sets of paths, and peg representations are compatible with locations.

We use linear-programming to maximize our objective function subject to these constraints. To do this, we introduce binary-weights, $\omega_{r,g}$, into our objective function, which determine if the pairing of a representation, r, to a geographic feature, g, will be included in the final tactile map. If the weight, $\omega_{r,g}$, is set to 1, then the pair, $p_{r,g}$, will be included in the optimized tactile map. Otherwise, the weight must be 0,

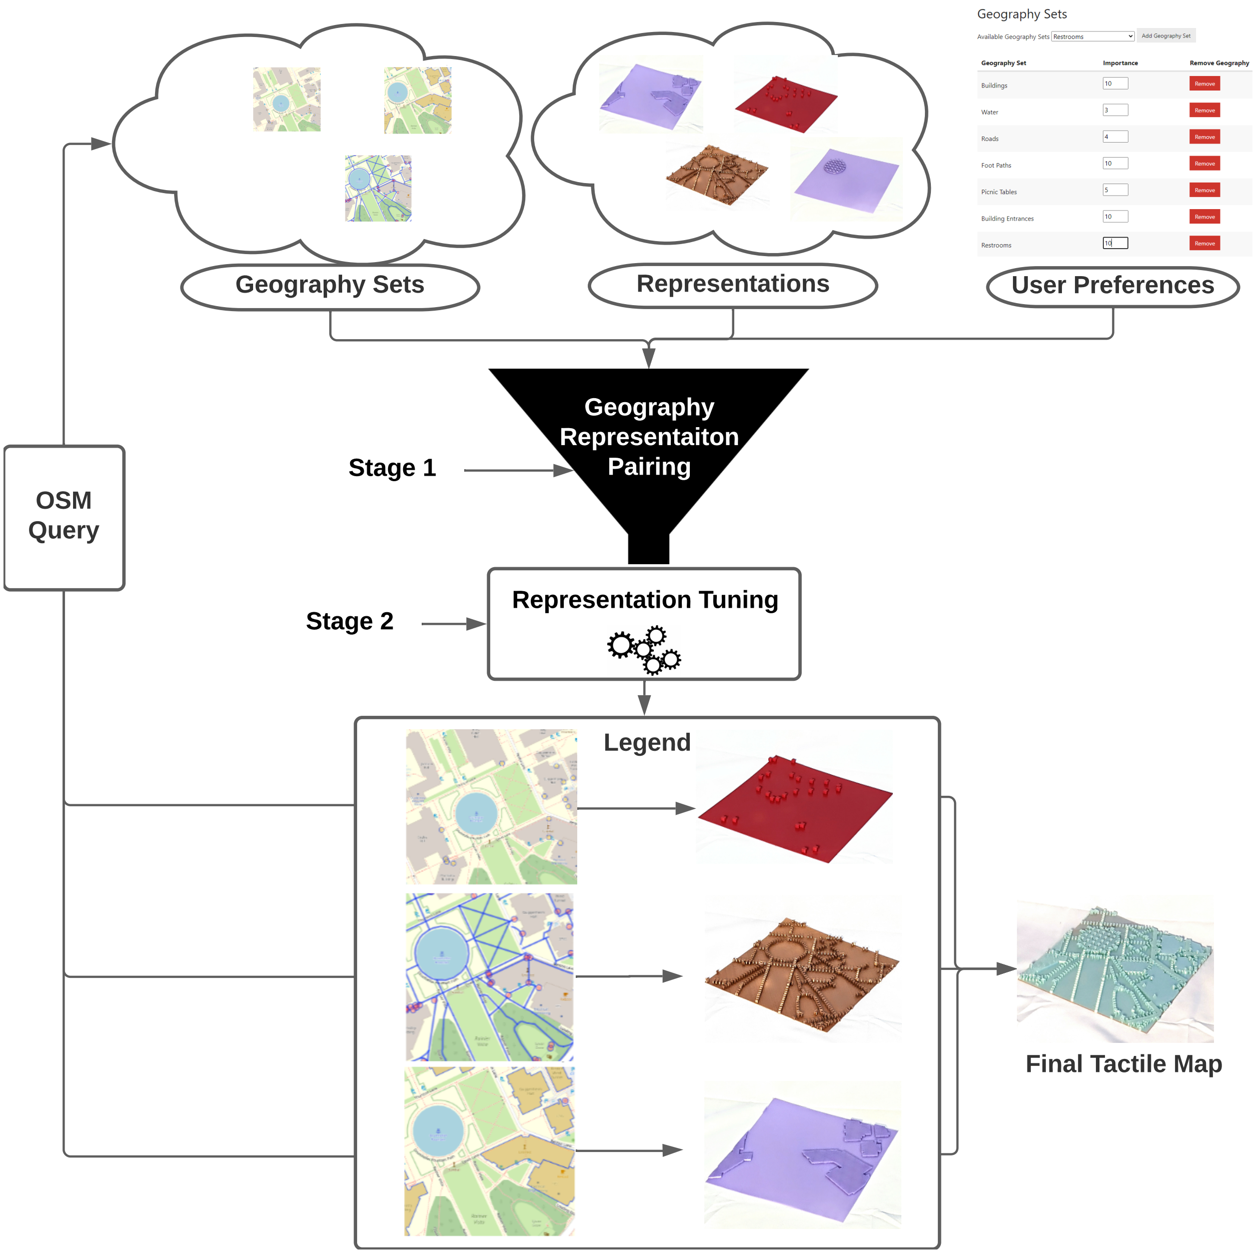

Figure 28: Maptimizer takes in a geographic feature, a representation set, and user preferences to pair representations to geographic features (stage 1). The representation parameters are tuned in stage 2. These representations take in geographic data from OpenStreetMap to combine representations into the final 3D printed map.

and the pair will not be included in the tactile map. The optimal set of pairings between representations and geographic features, $\mathbb{P}$, is a subset of all pairings, $\mathbb{P}_{all}$, of all representations, $\mathbb{R}$, to all geographic features, $\mathbb{G}$. To generate the set of optimal pairings, we reformulate our objective function and constraints as a solvable set of linear equations. The resulting weights are rounded to 0 or 1.

$$\text{Maximize} \sum_{p_{r,g} \in \mathbb{P}_{all}} \omega_{r,g} o(r,g)$$
$$\text{s.t.} \begin{cases} \forall_{r \in \mathbb{R}} \sum_{g \in \mathbb{G}} \omega_{r,g} \leqslant 1 \\ \forall_{g \in \mathbb{G}} \sum_{r \in \mathbb{R}} \omega_{r,g} \leqslant 1 \\ \forall_{p_{r,g} \in \mathbb{P}_{all}} \omega_{r,g} > 0 \implies r \text{ is compatible with } g \end{cases} \tag{9}$$

We formulate our pairing constraints as linear equations that isolate each pairing weight, and ensure representations and geographic features are not overused (Equation 9). First, we constrain the weights to ensure that no representation is used by more than one geographic feature. That is, for all representations $r \in \mathbb{R}$, the sum of the pairing weight between r and each geographic feature, $g \in \mathbb{G}$, is less than or equal to one. This implies that only one or none of the pairs will be used. Similarly, we constrain the weights to ensure that no geographic feature is paired to more than one representation. That is, for all features, the sum of the pairing weight between the feature and each representation is less than or equal to one. Finally, for each pair between a representation and geographic feature, the corresponding weight can only be greater than zero if the representation and feature are compatible.

9.2.5.2 *Stage 2: Representation Parameter Tuning*

The first stage of optimization will determine what geographic information will be included in the map and what types of representation will be used to represent it. However, we still need to tune the sizing parameters of each representation to maximize our objective function (e.g., adjusting the width of a path). Increasing a representation's parameters will increase its communicability (Equation 6). Conversely, decreasing these parameters will decrease the attention cost (Equation 8). By making these adjustments we are making trade offs between *communicability* and *attention cost* in our objective function (Equation 5a).

We implement this optimizer with OPTIMUM. First, we create objectives that measure communicability and attention-cost for each representation and geographic-pair. Next, we create modifiers that will effect these objectives. For each representation we provide four modifiers that increase/decrease the depth and width of the representation. The optimizer uses a design selector that selects, with a bias towards tactile map designs and towards those that have been most visited. This bias increases from 0.1 with each iteration. The modifier selector sorts modifiers by their history of improving the objective-score with a static 0.75 selection-bias. This OPTIMUM optimizer is configured to do Ant-Colony Optimization [58], and explores the search space of tactile maps to quickly converge on quality maps[1].

We chose Ant Colony Optimization after testing a variety of metaheuristic optimization methods; generally these methods find quality solutions in poorly defined and discrete search spaces, like the space of possible tactile-maps [32]. Ant Colony Optimization has had success in a variety of routing and scheduling problems [58], which have similarities to tactile-map optimization. We expect this method performs well because it "*intensely*" [32] searches a space of similar designs, before jumping to other portions of the design space. Since a high-quality starting position is produced by stage 1, an intense search is preferable to a diversified search that more readily jumps around the search space.

9.3 USER EVALUATION

We conducted a two part user-evaluation of the optimized tactile maps to answer two research questions. First, does the ability to customize what information is presented in a tactile map improve the experience

1 A more in-depth explanation of the Ant-Colony Optimization method is presented in the publication of this chapter [108]

for users? Second, does the combination of optimization with customization further improve the user's experience? Throughout our study, we considered three tactile map conditions: 1) *standardized-maps* made with TouchMapper [126] which cannot be customized by the user, 2) *customized-maps* that only include the most important information defined by a user, and 3) *optimized-maps* which consider a user's preferences, as well as other factors and are generated with Maptimizer. Participants were unaware of how any map was created until after they completed the study. We only told them that we had made these tactile maps and wanted their feedback. Given tactile maps of the same locations, we measure the quality of a tactile map based on a user's preference for different types of maps, and the user's ability to quickly and accurately find a specific location on a map using a verbal description. To test the quality of our *optimized-maps* compared to *customized-maps* and *standardized-maps*, we conducted a user study with six BLV participants from Seattle.

First, participants completed an online survey of their demographics and experience with tactile maps. The survey also included questions copied from the Maptimizer GUI, which helped us populate the interface fields when creating each participant's customized map. Example questions included the participant's preference on a scale from 0-10 for different textures and geographic features. The results of these surveys were used to generate each participant's customized and optimized maps.

Following the survey, participants then met researchers at a public location for an hour-long session where they used these tactile maps. Participants received a travel stipend to commute to the location and were compensated $40 for their time. Participants were recruited from the Seattle metro area, and some were familiar with the mapped locations, however, they were not told the locations until after the study was completed. In response to an open-ended question, two participants identified as men and four identified as women. Participant's ages ranged from 28 to 72 years of age (mean of 47 years, standard deviation of 18 years). Four participants identified as blind and two as having low-vision. On a scale from 1 to 5, 5 being very familiar, on average, participants rated their familiarity with tactile maps as 2.66 with a standard deviation of 0.5. Only one participant had no prior experience with tactile maps. We found no significant effects of participant's demographic information on the results of the study.

9.3.1 *Methods*

We conducted the think-aloud study in two parts. The first part helped participants familiarize themselves with the tactile maps and practice describing the representations to researchers. In the second part of the study, we tested whether maps that were customized or optimized affected participants' ability to identify a location on the map from a verbal description.

9.3.1.1 *Part 1: Familiarity with Tactile Maps*

Part 1 of the study primarily serves to help the participants adjust to the study protocol and familiarize themselves with the different representations in the tactile maps. Additionally, it allows us to compare how different ways of designing maps affect participant's preferences.

We first showed participants three different tactile maps of the same location in Seattle: a park that has roads, walkways, buildings, a lake, and some small locations of interest (e.g., picnic tables, benches, sculptures). This location was the simplest location in our study, serving as a simple case to help participants familiarize themselves with the different types of maps and representations. Three maps were created for our three conditions: *standardized-maps*, *customized-maps*, and *optimized-maps*.

STANDARDIZED-MAP (Figure 29a): Maps generated with TouchMapper [126] cannot be customized; all maps made with TouchMapper pair each geographic features to standard, unique representations. We use TouchMapper as a control condition across users.

CUSTOMIZED-MAP (Figure 29b): The customized maps consider the participant's preferences but do not use optimization methods. Each customized map uses the six most highly-ranked geographic features from the participant's preliminary survey. During pilot tests, we found that any more than six features consistently produced cluttered maps that were too difficult to read. These were each mapped to the most

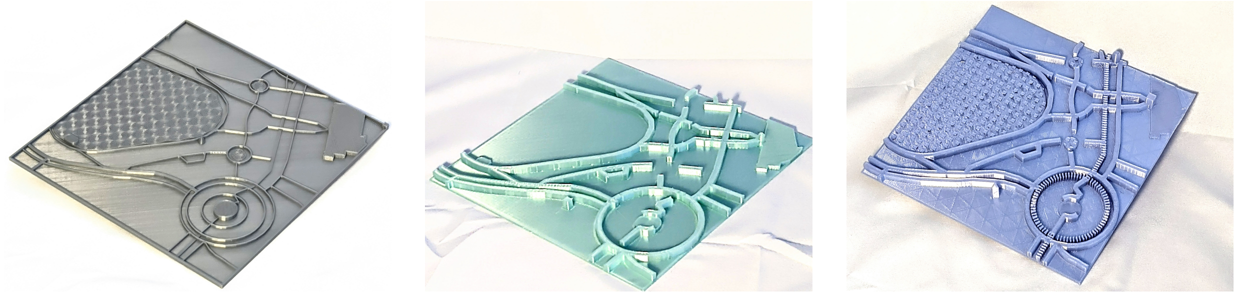

(a) Standardized Map (b) Customized Map (c) Optimized Map

Figure 29: Three maps of the same public park generated with: (a) a standardized map, (b) a user-customized map, and (c) an optimized map produced with Maptimizer.

highly-ranked, compatible representations available in Maptimizer. Standard depth and width parameters were used, rather than optimizing them. This experimental-condition helps us compare the effect of customization through the user's preferences without optimization.

OPTIMIZED-MAP (Figure 29c): By feeding the participant preferences from the preliminary survey into the Maptimizer interface, we generated optimized maps that considered both the participant's preferences and contextual information, such as the size of different geographic features of the location. We opted to use the preliminary survey, rather than having participants enter values directly into the interface during the study, because the tactile maps take at least 4 hours to print, and we needed to minimize participant's in-person interaction with researchers during the COVID-19 pandemic. This is our experimental-condition that helps us examine the effects of optimization.

We randomized the order in which these three maps were shown to the participants. Following a think aloud protocol, participants examined the maps and were free to ask questions about what different elements on the map represented. For instance, a participant could identify a rounded peg and ask what information it represented. We asked participants to verbally describe the elements of the map they wanted to know about, rather than allowing them to physically point out the feature. This think aloud approach gives us a better understanding of which representations were easiest to identify and distinguish from one another. For example, a participant could not just point out a round peg, they had to describe it as a peg (or equivalent phrase) and differentiate its shape from other pegs (e.g., squares and triangles). As participants explored the maps, they provided feedback around what they found useful or confusing, with respect to both the information conveyed and the textures used to represent them. After examining all three maps, we asked them which map they preferred and why.

9.3.1.2 *Part 2: Location Identification*

Next, we wanted to examine how customization and optimization effected the participants' ability to identify a location on a map, given a verbal description. We presented users with three new maps, but this time, each map was of a different location in Seattle—a university campus, a different park with a lake, a public market. We counterbalanced map-generation conditions with these locations across three groups to ensure that, if one location was more difficult to understand, this would effect each map-generation condition equally. Each group received the locations in the same order: campus, park, market. Two participants were randomly assigned to each group.

For each location, we read aloud a short verbal description of a location on the map. Features that were available in each map-condition could be used to identify this location. In the customized and optimized conditions, these features may not be present, depending on whether the user happened to have a strong preference for that type of location—a sculpture on a university campus, the end of a foot path leading to a beach in a park, or an information booth in the market. Each verbal description described the location as relative to an area (e.g., buildings, water), a path (e.g., road, footpaths), and a set of locations (e.g., benches). We read the description to the participants then handed them the map to start exploring. We would repeat the description as often as the participant requested, but did not provide any other details about the location. Just as in the first task, participants could ask what a feature of the map represented by describing it verbally.

Participants could give up their search for the location, or would announce when they believed they had found it. We recorded the time it took them to complete the task, whether or not they found the correct location, and how confident they were that they had found the correct location on a scale from 0 to 5; 0 was reserved for participants who did not identify a location, and 5 indicated the highest confidence in the location they found.

9.3.2 *Results*

Table 14: Summary of key statistics from the location finding task. The Maptimizer maps outperformed both alternatives on all measures.

	Standardized	**Customized**	**Optimized**
Preferred Map-Condition	2 (33%)	0 (0%)	4 (66%)
Identified correct location	2 (33%)	3 (50%)	6 (100%)
Confidence in location M (SD)	2.17 (2.04)	3.33 (2.25)	4.33 (1.03)
Time elapsed M (SD)	169.00 (112.93)	167.83 (99.72)	150.16 (111.54)

Based on the results (Table 14), we found that optimization helped participants identify a location. Overall, participant's preferences for different types of maps were dependent on how naturally specific representations paired with specific types of features. However, none of these generation methods consistently paired information this way. When we focused participants' efforts on using the maps to identify a location, participants had significantly more success with optimized maps. We suspect this is because the optimized maps provided both information the participants had ranked highly, and critical contextual information.

9.3.2.1 *Participant's Preferences*

The first part of the study showed that participants' preferences for different ways of generating a map varied. Overall, four people preferred the optimized map and two preferred the standardized map, while no one preferred the customized maps. While thinking aloud, our participants revealed insights into what made maps attractive. For instance, the standardized maps used a "*kind of rippled*" (P1) texture to represent water. Some participants (P1, P3, P6), intuitively knew this was water, but others assumed it was a "*hilly grass, kind of thing*" (P2). Alternatively, some participants based their preferences on how easy it was for them to identify information that was important to them after learning about the area. For example, P3 preferred the optimized map "*because it identifies the different items [she] needs*", such as water, the entrances to buildings, and a picnic table; whereas, in the customized map "*the water is not identified which would be a little confusing*". Notably, P3's customized map did not include water because she had ranked water features lower than other features. Since the water took up such a large portion of the map, Maptimizer added the water feature to provide context.

Most of our participants associated specific geographic features with representations that evoked an image (e.g., the rippled water texture). However, neither our customization condition nor our optimization method currently takes this into consideration. For instance, P2 did not think the square texture Maptimizer assigned to the lake matched his image of water, instead describing it as "*Arlington Cemetery for flees [sic]*" because the small rectangles appeared like very tiny graves. However, later during the location finding task, he found the information booth quickly because Maptimizer had assigned it a round peg representation. He said it reminded him of a lighthouse, "*what's the shining light coming out of this lighthouse? It's information*". Relating specific types of representations to geographic information is a difficult task, especially as different types of information are represented, creating conflicts. Exploring ways to include common representations of information in our optimization process is a promising area of future work.

9.3.2.2 *Effects on Location Identification*

While the number of participants is small, we found that Maptimizer's optimized maps helped participants identify locations more accurately. When presented with optimized maps, all six participants correctly identified the described location on their map, regardless of which group they were in. Accuracy varied when participants used standardized and customized maps. With the standardized maps, only the participants given the standardized park map could identify the correct location. Similarly, with the control maps, only half of the participants could identify the location. Two of these participants had the park map and the third had the market map. A χ^2 test shows that the map's location had a significant effect on success in location finding ($\chi^2 = 6.08, p < .05$). However, a χ^2 test demonstrates that the way a map was created (e.g., standardized, customized, optimized) also had a significant effect on participants' success at finding the correct location ($\chi^2 = 6.08, p < .05$). Since map location and the way a map was generated were counterbalanced, we suspect that these two effects are independent and the high success rate with optimized maps is attributed to the generation method.

9.3.3 *Evaluation Limitations*

There are two key limitations of this evaluation. The first, and most signification, limitation is the small number of participants. While our results are promising, with so few participants, we cannot know if Maptimizer would continue to perform well with new participants. Second, our study is limited to only a few locations. While we selected these areas to have a diverse set of features, we may not have captured the limitations of Maptimizer in other relevant contexts.

9.4 DISCUSSION

Maptimizer demonstrates how generative design for digital fabrication can bridge the gap between the Ability Based Design framework [269] and values of end-user empowerment [117]. Ability Based Design suggests that technology should adapt to users needs and abilities in an automatic way. However, unlike the digital user-interfaces this framework was designed for, the physical world is static and cannot be automatically adapted to changing contexts. 3D printing a tactile map takes time, digital fabrication takes time. However, this added time affords new opportunities to orthogonal experts in disability—the disabled users themselves. Design time is built into the digital fabrication process, and this empowers users to explicitly define their needs and abilities for the system.

Consider the principle of *performance* from Ability Based Design. Performance is the system's prediction of a user's performance within an interface. When an optimization is being run on dynamic software, every few seconds the system must model user's performance, rather than asking them to evaluate it. When digitally fabricating a design, the user can define a model of their own performance. By ranking geographic features and representations, users provide Maptimizer a model of how they will use the tactile map.

Alternatively, consider the principle of *context*, i.e., proactive sensing to anticipate changes in a user's abilities. The context of how a static tactile-graphic will be used is critical to identifying what information and representations should be prioritized and clearly presented in the final design. As in Maptimizer, end-users should be the ultimate authority on how the graphic will be used in a given context, because a generalized system cannot account for each user's unique needs.

9.5 CONCLUSION

Tactile Maps are useful tools to support navigation by people who are BLV. However, to make tactile maps more readily available, we need tools that can generate those maps for new locations without the support of a sighted cartographer. Beyond increasing availability, these tactile maps should be customized to meet the specific needs of an individual user. The design of each map should consider what information is most valuable to a particular person, and how information can be represented most effectively. However, requiring the user to fully define their tactile map creates a cumbersome design process which may not be worth their efforts. Instead, optimization techniques enable us to automatically adapt designs to user's specific abilities and needs. While Maptimizer is an assistive, rather than clinical, application, it demonstrates principles of clinical CAD tools and how meta-tools like OPTIMUM support their implementation.

10 KNITGIST: GENERATING FUNCTIONAL MACHINE KNITTED TEXTURES

While this thesis is focused on supporting digital fabrication for medial and assistive applications, the principles behind clinical CAD tools and the meta-tools that support their implementation can generalize to other burgeoning domains. This chapter describes a contemporary thread of research on automatic machine knitting and how OPTIMUM supports the implementation of a knitting generative design tool based on the orthogonal expertise of hand knitters.

Despite industrial knitting machines' wide adoption across the garment industry, there has been so little work on knitting-design tools, that an entire platform must be built from scratch. Recent advances in this area have almost exclusively focused on knitting 3D meshes [185], with little attention payed to the physical and aesthetic effects of knitted textures [131, 186]. Knitted textures produce the complex effects needed to knit garments [132] and novel soft mechanisms [11, 12, 160, 161].

The complex physical and visual properties of knit *texture* come from the combinations of numerous *stitches*. There are multiple properties of knitted textures that must be maintained to create a knit object that will not unravel. Maintaining these properties while adding requirements to generate specific textures constrains the problem further. Such design work currently requires extensive orthogonal expertise from knit designers. Rather than tediously specifying textures stitch by stitch, a design tool should enable designers to specify high-level objectives. However, generative design of knitted textures poses two key challenges; (1) maintenance of the hard constraints that ensure a knit object will not unravel [185], and (2) generation of user-defined physical and aesthetic objectives.

Consider the following scenario; a designer may wish to create a curly, symmetrical texture, without tediously specifying that texture stitch by stitch. Ideally, knitting domain-knowledge is already embedded in a generative tool using OPTIMUM. Using the heuristic-mapping between between objectives and modifiers, the designer can build an optimizer that creates this curly texture. First, they map an objective that the texture curls up on itself to a modifier that causes random stitches to curl forward. Then, they map two objectives; that the texture is vertically and horizontally symmetrical to two modifiers that, respectively, mirror stitches across the axes of symmetry. An experienced knitter would identify *stockinette* (a texture of all forward-curling stitches) as a simple and common texture that meets these criteria. The resulting optimizer consistently produces a stockinette texture. When the designer changes the objectives to target a variety of physical and aesthetic properties, the optimizer reliably produces other canonical textures (Figure 30).

10.1 BACKGROUND: AUTOMATIC MACHINE KNITTING

There are a wide range of rudimentary design tools for knitting. Commercial tools often target either hand or machine knitters, not both. Machine knitting is largely done with stitch-level machine knit charting tools and higher level 'wizard' tools, that limit designs [224, 234, 240] to a narrow set of garment templates. Hand-knitters rely on sourcing designs from books and repositories [33, 136, 208], or designing textures with adopted tools like spreadsheets [183]. These improvised methods do not support any verification of the final knit object. Ultimately, texture is reserved for hand-knitting experts.

Researchers have approached sizing as a challenge of manipulating 3D meshes that are hand [181, 271] and machine [185, 186, 202] knittable. The addition of texture, when supported [186], may change the shape in unpredictable ways. Simulation of knitted fabric has shown promising results [47, 128, 176], including supporting interactive design of small texture patches [148], but still requires case-specific hand-tuning to provide results specific to a given yarn.

Recent work has focused on algorithmic, verifiable, computational solutions for generating knitting patterns. McCann *et al* created a machine-knitting compiler, which included a simple, machine-level language for controlling knitting machines [170]. The simplified machine language, Knitout [169], is an instruction

(a) Stockinette: unrolled and rolled

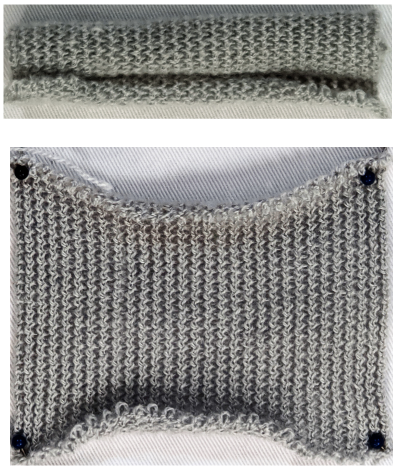

(b) Cables: unrolled and rolled

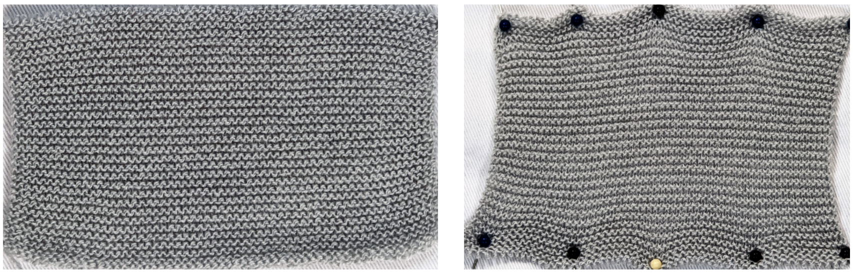

(c) Garter: un-stretched (left) and stretched (right)

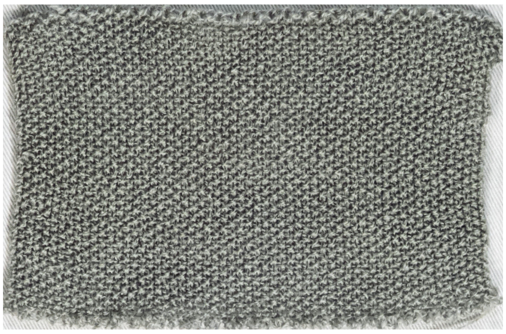

(d) Seed-Stitch: horizontally and vertically shrunken

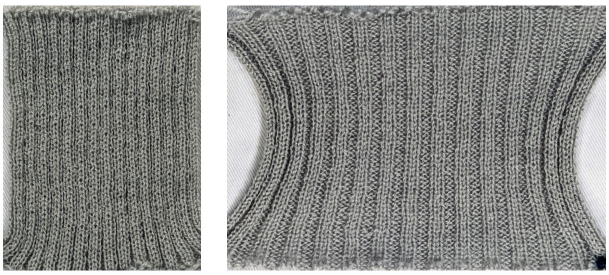

(e) Knit 2 Purl 2 Rib un-stretched (left) and stretched (right)

Figure 30: Exemplar textures produced during the described scenario that are curly (a), cabled (b), vertically-elastic (c), shrunken (d), and horizontally-elastic (e).

set for assigning loops of yarn to be held on beds of hook-shaped needles[1]. Under this instruction set, Transfer Planning is used as the process of assigning needle locations to loops in a graph structure, such that the represented knit object can be knitted on a knitting machine [154].

Overall, algorithmic machine knitting has formed a cohesive architecture and work flow: (1) model knit object by 3d modeling [186] or texture programming [131], (2) convert (e.g., compile) those models to KnitGraphs to be verified, manipulated, and evaluated, and (3) convert the KnitGraph to knitting machine instructions by transfer planning [154]. While there has been more substantial work on generating KnitGraphs from 3D modeled shapes [185, 186], little work has focused on generating functional texture programs based on user specifications and assurance of knittability.

10.1.1 *The Structure of Knitted Texture and Objects*

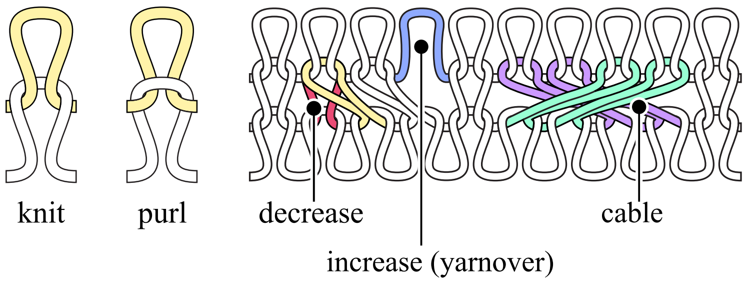

Figure 31: Loop-to-loop structures of common stitches: knits, purls, decreases, increases, and cables

A knitted fabric starts with a row of loops, and then additional rows of loops are created by pulling loops through the top-most row to create a sheet of fabric. A single loop is not stable; pull on its ends and the loop falls apart. Knitted fabric gains its stability from the relationship between loops; when pulling a child loop through a parent loop, the parent loop becomes stable. A loop pulled through another loop is called a *stitch*, adjacent stitches are called *rows*, and columns of stacked stitches are called a *wale*.

Texture derives from interconnected stitches. There are three composable properties of a stitch that change its effect on the texture: (1) the orientation a loop is pulled through another loop, (2) how many loops it is pulled through, and (3) how many other loops it crosses over. Hand knitters have developed a large set of named stitch-types that cover a variety of the most useful combinations: knits and purls, increases and decreases, and twists and cables.

KNIT/PURL Loops can be pulled through other loops in either of two directions. The most basic stitch is a *knit*—a loop pulled from the back of the fabric, through another loop, to the front (Figure 31). A *purl* is the opposite: a loop pulled front-to-back through another loop (Figure 31). A single stitch in isolation cannot meaningfully be labeled as one or the other: a purl is simply a knit viewed from the back.

The distinction between knit and purl becomes important at the textural level, with multiple stitches involved. Each stitch has a tiny amount of springy curvature, due to its out-of-plane structure, and this curl accumulates across collections of stitches. A knit fabric in which all stitches are oriented the same way will tend to roll up; textures which vary the stitch direction, such as "rib" (alternating columns of knits and purls) or "garter/welting" (alternating rows of knit and purl), exhibit other behaviors, such as scrunching or puckering.

DECREASE/INCREASE More than one loop can be pulled through another loop (an *increase*), and a loop can be pulled through multiple other loops (a *decrease*). For textures, decreases and increases are locally paired (canceling each other out) to produce lace patterns. For example, a yarn-over leaves a small hole in the fabric when paired with a decrease. Special types of increases and decreases are used create the first loops in a knitted object and to stabilize the last loops in the object. Cast-ons are increases that increase the number of loops on the first row so that they can be pulled through subsequent rows. Bind-offs are

1 For additional knitting machine details, see [11, 170]

decreases on the last row that decrease loops on the same row, until only one loop is available, which is tied off.

TWIST/CABLE A stitch can cross over neighboring stitches. *Cables* are formed by transposing adjacent sets of stitches in the same row. A cable of two stitches is called a *twist*. Cables tend to stiffen the fabric by creating additional tension on the loops as they are stretched across other loops.

10.2 REPRESENTING KNITTED TEXTURES AS GRAPHS

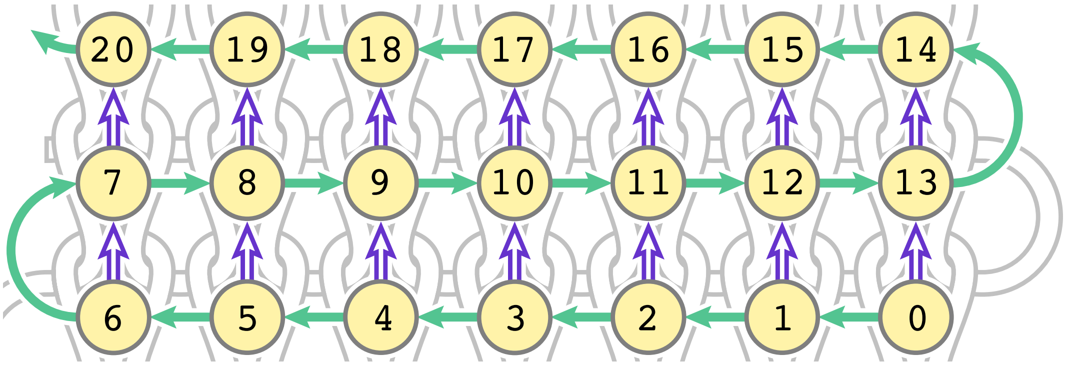

Figure 32: A KnitGraph is a link graph representation of a knitted fabric. Loops (yellow-circles) are connected by yarn edges (green) and stitch-edges (purple).

Table 15: Summary of KnitGraph notation.

Notation	Interpretation
$u < l$	Loop u was constructed before l
$u \to l$	Loop u comes just before loop l on the yarn
$p \uparrow l$	Loop p is pulled through loop l, creating a stitch edge
$o(p \uparrow l)$	The orientation of the stitch edge $p \uparrow l$
$P(l)$	Set of parent loops l is pulled through
R_i	The row at index i
N_R	Top row-index in a texture
$i(l)$	The index of l in its row

Knitted objects are represented as *KnitGraphs* that can be automatically evaluated and manipulated while maintaining desirable properties. A KnitGraph consists of a set of *loops*, with directed *stitch edges* representing where a loop is pulled through a parent loop. A loop may have multiple or no *parent* loops, but can only have one child loop. Each stitch edge has an *orientation*, denoting if the loop was pulled back-to-front (knit) or front-to-back (purl). Stitch edges may cross over one another in a cable stitch. Each loop is made with a yarn. Multiple yarns can be used to create a knitted object, and alternating between different colored yarn can create images and colored patterns in the texture.

KnitGraphs are directed graphs where nodes represent loops with yarn-wise and loop-to-loop edges (see Figure 32). A KnitGraph is: an ordered set of loops on a yarn[2], $l \in Y$; a set of yarn-wise edges, $n \to l \in E_Y$,

2 Generally, when referring to loops I use the terms l, m, p, and n. Generally, l is an arbitrary loop, p is l's parent, m is l's child, and n is l's yarn-wise neighbor.

between loops in the order they are constructed; and a set of loop-to-loop edges, $p \uparrow l \in E_L$, representing how loops are pulled through other loops. Each loop-to-loop edge is labeled with an orientation: a loop pulled back-to-front or pulled front-to-back. By convention, the first yarn edge is directed from right-to-left when the cloth is viewed from the front. A summary of KnitGraph notation is provided in Table 15.

A stable KnitGraph has two properties that guarantee that it will not unravel and that it can be constructed on a knitting machine.

PROPERTY 1: LOOP-TO-LOOP STABILITY The primary constraint of a knitted object is that each loop must have at least one other loop pulled through it. Any KnitGraph that satisfies this property will not unravel. For every loop p there exists a loop l that is pulled through p.

$$\forall p \in Y \; \exists l \in Y : p \uparrow l \in E_L \tag{10}$$

PROPERTY 2: TIME ALIGNED LOOPS During knitting, yarn-wise edges establish the relative horizontal position of neighboring loops and the time that they are constructed. If a loop, l, is pulled through a parent loop, p, p must be constructed before l.

$$\forall p \uparrow l \in E_L : p < l \tag{11}$$

10.2.1 *Generating Machine Knitting Instructions*

Machine knitting instructions that will fabricate the represented knit-object can be derived from KnitGraphs. Knitting machines use rows ("beds") of hook-shaped needles [3]. In machine knitting, each loop in the last row is held on its own needle, and loops can be transferred from needle to needle. Unlike in hand knitting, in which the most recent row of loops is free to slide along the single long needle, each column of machine-made stitches is held at the top loop by its own separate needle. Thus, each loop must be allocated a specific needle at the time of its construction, and its parent must be located there at that time to receive it. Combining loops onto needles for decreases, creating spaces for increases on empty needles between loops, and using the front and back beds for knitting and purling can all require re-arranging loops between rows of knitting.

By convention, the first row is allocated right-to-left, with each loop assigned a needle directly leftward of the one before it. For subsequent rows, two variables are maintained: a *cursor*, corresponding to the loop's position in the row, and a *slide* variable. While iterating over the loops from left to right, *cursor* is incremented and *slide* is updated per loop: if a loop has one parent, *slide* remains the same; if a loop has more than one parent, *slide* is decreased for each parent; if a loop has no parents or its parents have other children, *slide* is increased. Each loop is allocated to a needle at position = [cursor] + [slide]. The allocated needles are then used to determine the amount that each parent will be offset to support the new row of loops. Once the needles are allocated, each loop's parent(s) might need to be moved to the location where that new loop will be formed. These re-arrangements are accomplished via needle transfers determined by a "schoolbus+sliders" transfer solver [154].

10.3 KNITGIST HEURISTIC LIBRARY

Manually designing knitted textures to have specific properties is a difficult and tedious task. *KnitGIST* is a generative design tool implemented with OPTIMUM that modifies KnitGraphs to match a designer's aesthetic and functional goals. KnitGIST includes a library of eight parameterizable objectives and five modifiers that support the generation of optimized knitted textures (Table 16). Heuristic maps relating these objectives and modifiers can be used with OPTIMUM's default simulated annealing optimizer to generate a variety of complex knitted textures.

3 See [11, 170]

Table 16: KnitGIST Library broken into Objectives and modifiers.

Component	Function	Parameters
Objectives	Curl	NA
	Shrinkage	Direction: Horizontal or Vertical
	Elasticity	Direction: Horizontal or Vertical
	Opacity	NA
	Symmetry	Axis Location
	Style	Style-Type: Knit-Purl, Cable, Lace
	Color-Match	Desired Color
	Imagery	Region-Map, objectives, objective-Weights
modifiers	Flip Stitch	NA
	Lean Stitch	NA
	Replace Stitch	New Stitch Type
	Mirror Stitch	Axis Location
	Swap Stitch	Alternate KnitSpeak-Texture

10.3.1 *Objectives*

Objectives evaluate the properties of the knitted texture. The KnitGIST library includes four physical and four aesthetic objectives (Table 16). The objectives were driven by prominent categories laid out by hand knitters [236]. In the KnitGIST library, each objective scores a texture by comparing a value of interest, v, to a designer-defined target, t, over some bounds, $[min, max]$. The score is one if the calculated value is equal to the target, and drops off to zero linearly as the value approaches the bounds Equation 12. Each objective calculates v differently.

$$S(t, min, max, v) = \begin{cases} 0 & v \notin [min, max] \\ 1 - \frac{t-v}{t-min} & v \leqslant t \\ 1 - \frac{v-t}{max-t} & t < v \end{cases} \tag{12}$$

The physical objectives presented describe how the resulting texture will modify a physical characteristic of textures: curl, shrinkage, elasticity, and opacity. How much a knit object curls, shrinks, stretches, or blocks light is ultimately dependent on a wide range of manufacturing variables (e.g., machine, machine settings, fiber, gauge). However, texture plays a critical role in modifying the basic properties. For instance, an inelastic fiber can still result in a stretchy fabric if it is knitted into a horizontally-elastic rib texture. The following physical objectives produce a value, v, between -1 and 1, where zero implies that a texture contributes nothing to this specific property of a knit object. One or negative one implies that the texture maximizes this property of a given knit object. The negative scale implies direction, in the cases of curl (curl forward or backwards), and elasticity, where a texture may resist stretch.

10.3.1.1 *Curl*

To estimate the curl of a texture, v_c, it is important to understand why knitted textures can curl. Curl is an effect that occurs only in the vertical direction, due to the orientation of stitch edges between child and parent loops. This effect is cumulative. If the curl of stitches in a wale are all in the same direction, the last stitch in the wale is pulled forward causing the texture to curl. When stitches do not curl in the same direction, they cancel each other out locally. The effect of the curl of each successive stitch decreases with distance.

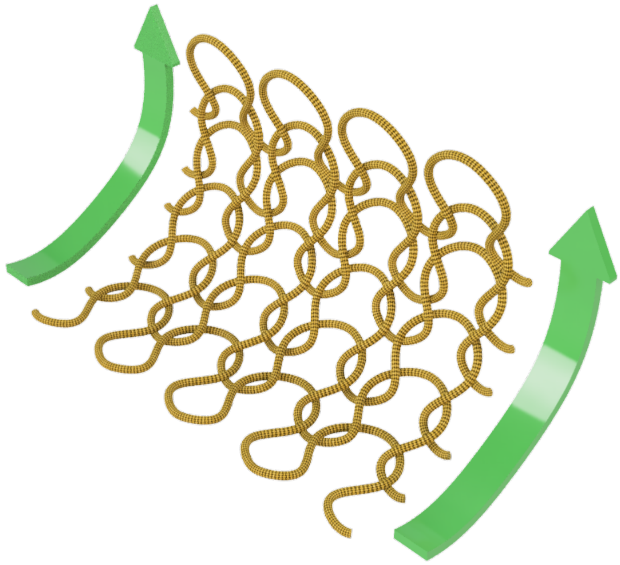

(a) All-Knits

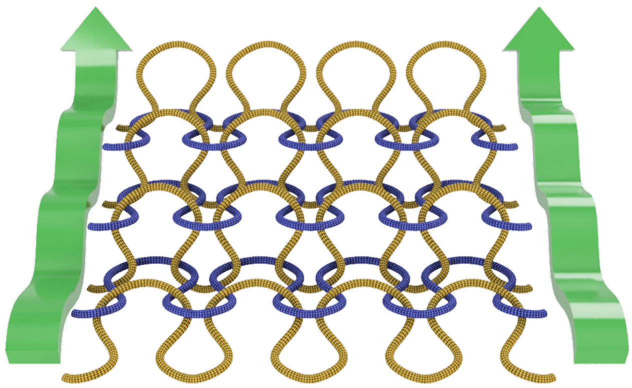

(b) Rows of Knits then Purls

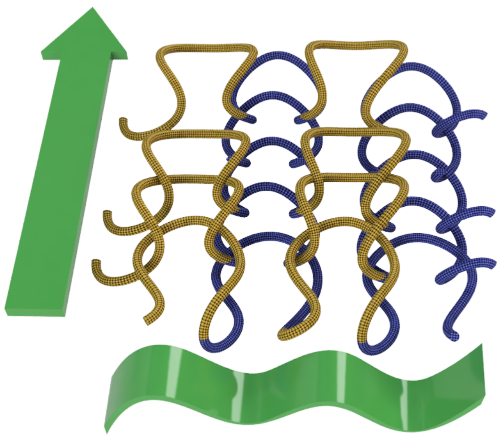

(c) Wales of Knits then Purls

Figure 33: Knits (yellow) tend to pull the fabric forward, while purls (blue) tend to pull the fabric back. A fabric with many connected knits will tend to curl (a). Switching between knits and purls will contract that curl, causing shrinkage (b and c).

Curl is modeled recursively (Equation 13a). If a loop, l, has no child, its curl, $c(l)$, is the average orientation across its parent loops. Otherwise, its curl is that average plus the cumulative curl of its child loop, m, divided by a decay-factor, σ. Empirically, two is a good value for σ. We estimate the cumulative curl of a knitted texture, ν_c, as the average curl of loops in the bottom row, R_0, which accumulates the curl of all subsequent rows Equation 13b.

$$c(l) = \frac{1}{|P(l)|} \sum_{p \in P(l)} o(p, l) + \begin{cases} \frac{1}{\sigma} c(m) & \exists m | l \uparrow m \\ 0 & \text{otherwise} \end{cases} \tag{13a}$$

$$\nu_c = \frac{1}{|R_0|} \sum_{l \in R_0} c(l) \tag{13b}$$

If a fabric is made up of only knits or only purls it will curl (Figure 33a). However, alternating patterns of knits and purls will prevent the fabric from curling. Curls may cancel out; for example, the forward curl of a row of knits is canceled out by the backwards curl of a row of purls (Figure 33b). Also, the forward curl of a wale of knits is canceled out by the backwards curl of a wale of purls (Figure 33c). Examples of these effects were demonstrated in the exemplar textures at the beginning of this chapter (see Figure 30).

10.3.1.2 *Shrinkage*

When a texture is constructed with alternating orientations of stitch edges, it can shrink significantly. Given an estimator for shrinkage, this objective can ensure that a texture is compact. Shrunken textures are frequently used at the boundaries of other textures or the edges of garments to prevent curl. Shrinkage is caused by the counteracting orientations of knits and purls, which causes stitches to overlap. A second cause of shrinkage occurs when loops are tightened by being pulled through a far away loop; this can be seen in cables, where loops are crossed over other loops, or in decreases, where loops are gathered through a child loop.

Shrinkage is caused by the relationship between neighboring loops in either a vertical (wale-wise) or horizontal (row-wise) direction. The wale wise neighbor of a loop, l, in a given direction, d, which is either vertical, $\updownarrow$, or horizontal, $\leftrightarrow$, is denoted as $\mathtt{neighbor}(l, d)$ (Equation 14). The loop n is a wale-wise neighbor to l if and only if n is pulled through l; that is, there is a stitch edge from l to n. Loops l and n are row-wise neighbors if, and only if, the row-wise index of n is one more than the row-wise index of l, and n and l are in the same row.

$$\mathtt{neighbor}(l, d) \iff \begin{cases} \exists l \uparrow n & d = \updownarrow \\ i(n) = i(l) + 1; n \in R(l) & d = \leftrightarrow \end{cases} \tag{14}$$

With respect to orientation, if two neighboring stitch edges have opposite orientations, they will tend to overlap each other. In other words, when a purl has knits on either side, only the knits are visible, because the knits on each side of the purl overlap it completely. This effect is visible in Figure 33c. This is true whether the stitches are aligned vertically (along a wale, in a parent-child relationship) or horizontally (along a row). Consider two stitch edges (between two loops each) where the child loops are either wale-wise or row-wise neighbors. Given loop l and its wale-wise or row-wise neighbor n, calculate the overlap, s_o, between two loops (l, n) as the difference between the average orientation of the stitch edges between the parents of l ($p \in P(l)$) and the parents of n ($m \in P(n)$), where $o(p \uparrow l)$ denotes the orientation that l is pulled through p. The term d denotes the direction (either horizontal or vertical) that determines which neighboring loop, n, is selected by the designer. So given loop l with the set of parent loops $P(l)$ and a neighboring loop n, with the set of parent loops $P(n)$, the loops overlap by:

$$s_o(l, d) = \frac{1}{|P(n)|} \sum_{m \in P(n)} o(m \uparrow n) - \frac{1}{|P(l)|} \sum_{p \in P(l)} o(p \uparrow l) \tag{15}$$

Loops overlap by the average of their *width*. The width of a loop depends on the location of its parent loop. When loops are pulled through a parent loop directly below them, they have a standard width. As the distance between the parent and child loops is increased (e.g., cable, decrease), the child loop is stretched vertically, making it thinner. As the distance between the loop l and its parent p increases, l is stretched thin. This distance is the difference between the in-row index of a loop, $i(l)$, and of its parent loop, $i(p)$. The width of the child loop, $w(l)$, is calculated as a factor of this sum of the distances between the loop l and all of its parents, $p \in P(l)$, plus the distance between the rows (i.e., 1) (Equation 16). The width of a loop, l, pulled through parents, $P(l)$, is:

$$w(l) = \frac{1}{1 + \sum\limits_{p \in P(l)} |i(l) - i(p)|} \tag{16}$$

We calculate the shrinkage between two loops as the average of their widths (Equation 16) multiplied by the amount they overlap (Equation 15). So for a loop l with neighboring loop n:

$$s(l, d) = \frac{(w(l) + w(n))}{2} s_o(l, d) \tag{17}$$

Shrinkage of a texture is the average shrinkage across all loops. A designer-set parameter d dictates whether to calculate shrinkage, either vertically or horizontally, by using the appropriate definitions of a neighbor (Equation 14). So over a whole set of rows of size N_R, we calculate the average shrinkage over all loops, l, in all rows as R_i.

$$v_s(d) = \frac{1}{N_R} \sum_{i=0}^{N_R} \left(\frac{1}{|R_i|} \sum_{l \in R_i} s(l, d) \right) \tag{18}$$

Consider a garter texture, made up of alternating rows of knits and purls. The rows will have opposing curl, so the stitches will overlap and the texture will shrink vertically. Alternatively, consider ribbing, made up of alternating wales of knits and purls, which shrinks horizontally.

10.3.1.3 *Elasticity*

Textures tend to be either vertically or horizontally elastic, as stretching in one direction shrinks the fabric in the other direction. Textures that do not shrink resist stretching. Similarly, textures that shrink in both directions resist stretching, because stretching the texture one way will cause stitches in the other direction to overlap. Consider again alternating rows of knits and purls and alternating wales. The alternating rows shrink vertically, but not horizontally, so the texture stretches vertically. Alternating wales do the opposite, so the texture stretches horizontally. Elasticity in one direction is the difference between the shrinkage in the target direction and the shrinkage in the opposite direction (Equation 19).

$$v_e(d) = \begin{cases} \max(v_s(\leftrightarrow) - v_s(\updownarrow), 0) & d = \leftrightarrow \\ \max(v_s(\updownarrow) - v_s(\leftrightarrow), 0) & d = \updownarrow \end{cases} \tag{19}$$

10.3.1.4 *Opacity*

Opacity is largely defined by pairs of increases and decreases, leaving a gap between them. Yarnovers are an increase where an extra loop, with no parent loops, adds one loop to a row. When this is combined with a decrease in the same row, it leaves an obvious hole or eyelet in the fabric. The opacity of the texture is the density of opaque loops (loops with a parent loop) and non-opaque loops (loops without a parent loop) (Equation 20). Intuitively, lace textures with many increase-decrease pairs will be less opaque than other textures.

$$v_o = \frac{1}{N}\sum_{i=0}^{N}\left(\frac{1}{|R_i|}\sum_{l\in R_i}\begin{cases}0 & |P(l)|>0\\ 1 & |P(l)|=0\end{cases}\right) \tag{20}$$

10.3.1.5 *Symmetry*

Symmetry creates aesthetically-balanced texture. We evaluate symmetry across the stitch edges equidistant across an axis in a texture. The designer chooses the axis location, defaulting to the center. Symmetry is a concept that can be specified by the designer. Generally, symmetry is binary, but there are many properties that may be symmetrical. By default, this objective compares three properties of the stitch edge: orientation, depth, and lean. Given a set of symmetry functions comparing two stitch edges, S, which return a value between 0 and 1, this framework calculates overall symmetry for a pair of edges by averaging the values returned by those functions. Over an entire texture, the symmetry value is the average symmetry value between paired equidistant stitches across a symmetry axis. Given a horizontal or vertical axis, the set A contains all pairs of edges equidistant from that axis. So a pair of edges $p \uparrow l$ and $p' \uparrow l'$ are in A if l and l' are equidistant from the axis, p is a parent of l, and p' is a parent of l'.

$$v_{sym}(A) = \frac{1}{|A|}\sum_{(p,l,p',l')\in A}\left(\frac{1}{|S|}\sum_{sym\in S} sym(p\uparrow l, p'\uparrow l')\right) \tag{21}$$

10.3.1.6 *Styles*

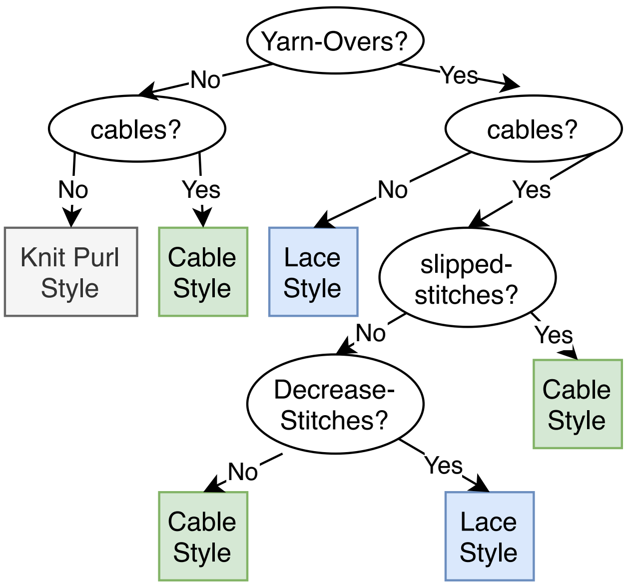

Figure 34: Style classification decision-tree model.

Styles are the broad categories of knitted textures used by most hand-knitters (i.e., knit-purl, cables, lace). Each of these styles have unique physical and aesthetic properties. Knit-purl patterns are the simplest, consisting of only knits and purls. Cable patterns use cable stitches to cross stitches over one another, giving the appearance of traveling-stitches. Lace patterns balance increases and decreases to create eyelets.

We estimate style with a decision tree, based on the types of stitches in a texture. We manually labeled 1548 knitting patterns from the Stitch-Maps repository [33] as *Knit-Purl*, *Cable* or *Lace* by examining swatch-photos. We calculated a variety of features from the KnitSpeak text including: stitch-counts, patterns of repeated combinations of stitches, and presences or absences of certain stitches in the texture. Next, we trained a decision tree to classify textures using 25% of the samples in a development set, 50% in a cross-validation set, and the remaining 25% as hold-outs for final testing. The model produced by the C4.5 decision tree algorithm (confidence interval .15, minimum of five instances per leaf) had an accuracy of 95.6% ($\kappa = 0.86$) over the withheld set. The best model used features specifying which stitch types were present (Figure 34).

10.3.1.7 *Color-Match*

The simplest aesthetic objective in the KnitGIST library compares the color of the KnitGraph, what colored yarn makes up the loops, to a desired color set by the designer. The objective returns the proportion of loops which match the desired color.

10.3.1.8 *Imagery*

Imagery is a way of creating visual effects such as a hexagon using texture. Imagery-objectives allow designers to apply other objectives over specific 2D regions of the textures. Designers specify a region-map where a set of loops are mapped to another objective (i.e., curl, shrinkage, elasticity, opacity, symmetry, style, color-match). Like all the previous objectives, the designer provides a target for the region's objective and acceptable bounds. The value of the objective is calculated over the regional subset of loops, rather than the whole texture.

Since it would be tedious to assign loops to a region-map by hand, we provide a simple painting tool to assign objectives to regions on a grid. Multiple imagery objectives can be defined over different regions with the same sub-objective and different target values. Target values between 0 and 1 are represented as the range of colors between blue and red. For example, as in Figure 35, the blue region will have a target value of 0 and the red region will have a target value of 1. The painting tool can be used multiple times when creating a heuristic map to create regions with different objectives. Previously painted regions can be loaded in, so the designer does not have to paint the same region repeatedly.

10.3.2 *Modifiers*

Each of the modifiers in the KnitGIST library changes a random stitch in six different ways. Note this is not an exhaustive set of possible stitch-level modifiers, however, it is effective at producing a wide variety of functional textures.

1. *Flip* its orientation (i.e., knit vs purl).
2. *Lean* it in the opposite direction (i.e., left vs right).
3. *Replace* a stitch with another stitch type.
4. *Swap* a stitch across a symmetry-axis
5. *Color* of the child loop
6. *Copy* a stitch from another texture that scores higher on the relevant objective

10.4 DEMONSTRATIONS

Combining KnitGIST's objectives and modifiers in different heuristic maps enables designers to create a variety of functional and attractive knit textures. Knit-designers and hand-knitters can recognize the relationships between stitch properties and different objectives to form heuristic maps.

10.4.1 *Optimizing Textures*

In the introduction, we demonstrated how a simple mapping between objectives and modifiers produced a set of canonical knitted textures that curl, shrink, and stretch. This demonstration builds on this scenario

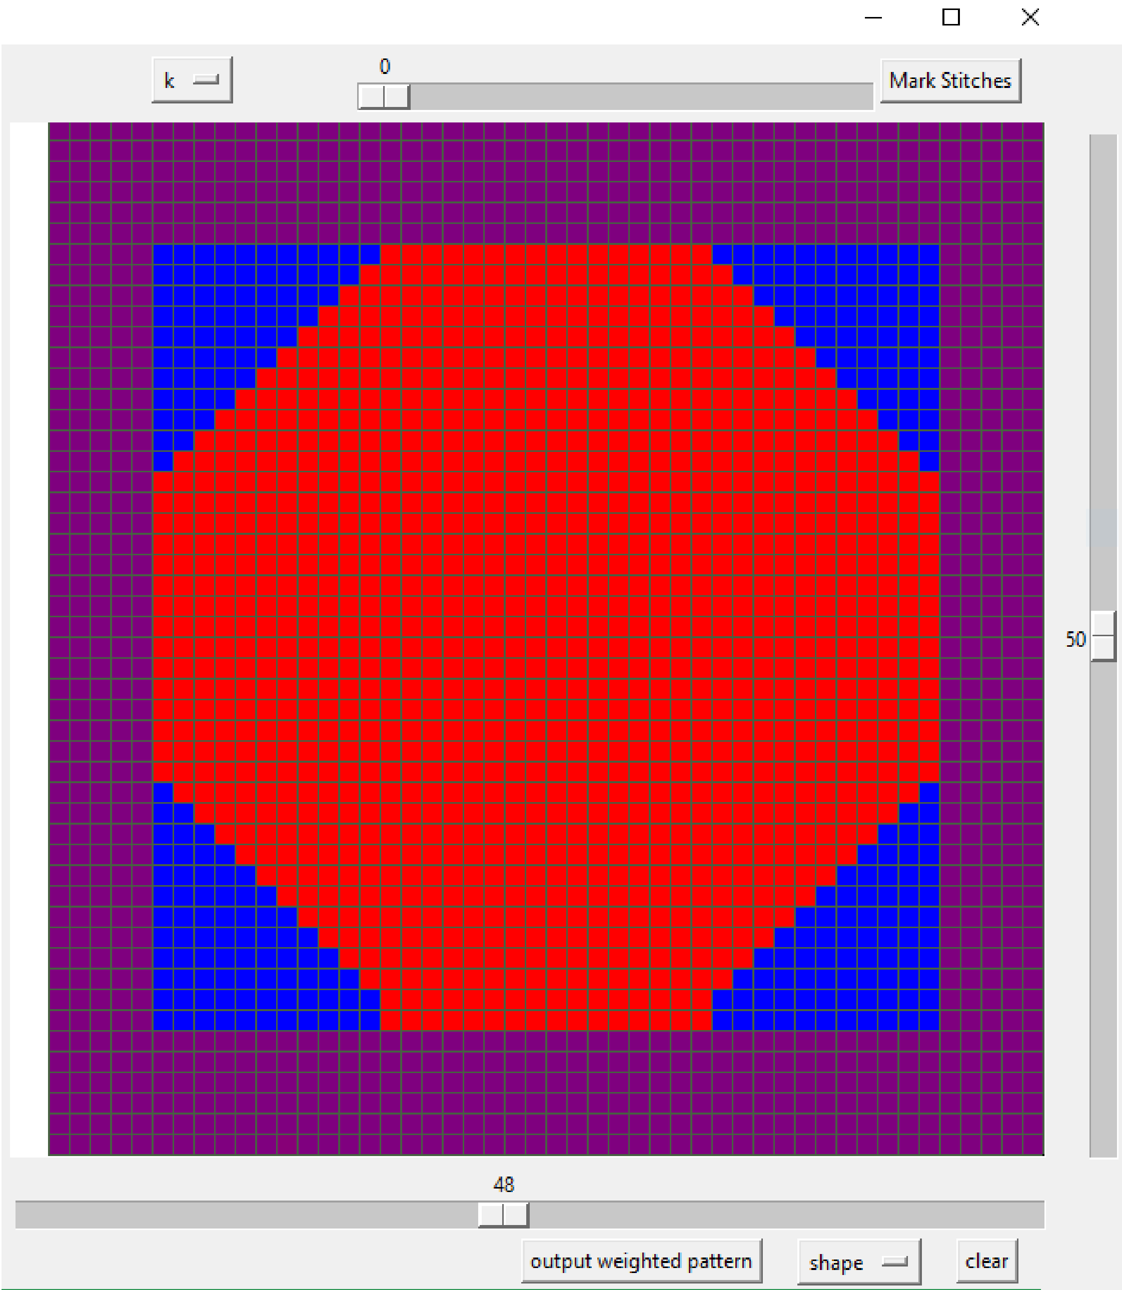

Figure 35: Example of imagery region mapping painting tool.

to produce a variety of welting textures with desired aesthetic and physical properties. Welts are textures with alternating horizontal stripes of stitches with opposite orientations. The constant orientation of each stripe causes the stripe to curl; switching to the opposite orientation causes the texture to curl backwards and shrink vertically. This produces a vertically elastic texture.

We use three objectives to generate a plain knit-purl welt. The first is a region map that marks the stripes that should curl forward. The second maps stripes that curl backwards. The third objective sets a high target vertical-elasticity. The curl objectives map to a flip modifier, which increases curl by aligning stitch orientations. The same modifier is mapped to the elasticity-objective. In addition to the flip modifier, we map a secondary mirror modifier, which replaces a stitch with a stitch equidistant across a center-vertical axis. This will increase the vertical symmetry of the texture. While the symmetry objective is not being used, vertically elastic textures tend to be symmetrical, so this approach is a useful heuristic.

(c) Basic Welt

(f) Welt with Lace

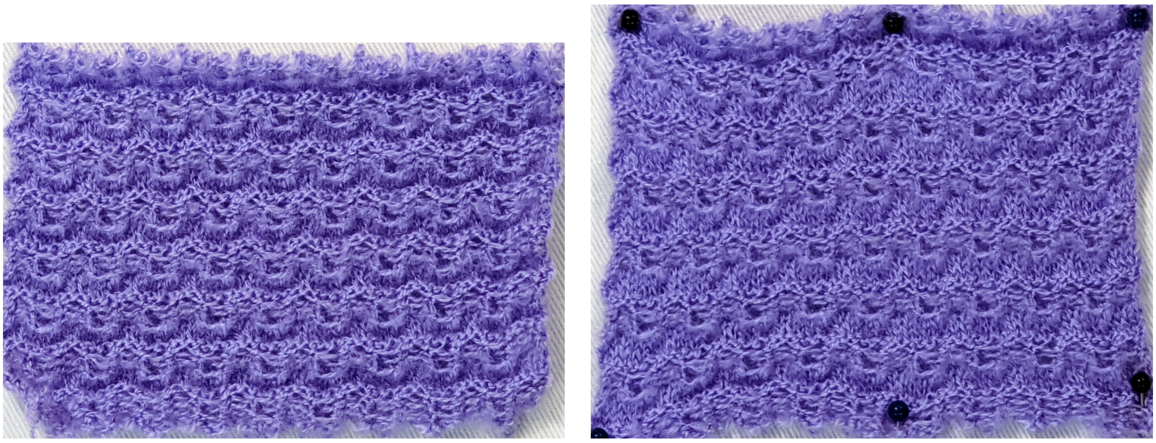

(i) Welt with Cables

Figure 36: The generated welt textures use functional objectives that alternate the direction of curls in horizontal bands. Aesthetic properties are applied to the Lace (b) and Cable (c) welts. The alternating curls maximize vertical elasticity.

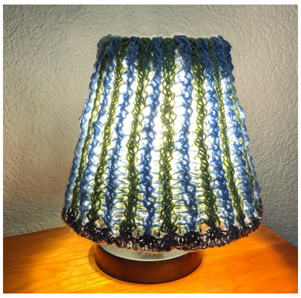

(a) Lamp Shade

(b) Crutch Cosy

(c) Color Work Pillow

Figure 37: KnitGIST supports the design of a variety of knit objects

Combining these objectives and trading out different styles (e.g., lace, cables) can produce a variety of effects (Figure 36). The heuristic map for each of these welts is described in Table 17.

Table 17: Welt Heuristic Map

Welt Type	Objective	Objective Weight	Modifier	Modifier Weight
All	Forward-Curl Stripped Region	3	Flip Stitch Orientation	1
All	Backward-Curl Stripped Region	3	Flip Stitch Orientation	1
All	Maximize ↕ Elasticity	2	Flip Stitch Orientation	2
			Mirror Stitch	1
Lace	Lace Style	1	Replace with increase-decrease pair	1
Cable	Cable Style	1	Replace with cable stitch	1

10.4.2 *Designing Knit Objects*

KnitGIST can also be used to design functional knit objects. Consider the properties of a lamp shade. It should allow some light to shine through, but diffuse enough to dampen the bulb. Additionally, lampshades are often conical, with a top that is narrower than the base. This can be done by shaping the knitted texture, but it can also be shaped by a rectangular swatch of fabric that is highly-horizontally elastic, which allows the top to remain compressed while the base is stretched wide. Finally, the top and bottom of the texture should bend or curl to cover the top and bottom of the lampshade frame. By bending around the frame, the only post-processing required will be to seam the texture into a tube and stretch it over the frame.

We created a heuristic map to specify these goals (Table 18) and generate a knitted texture that would meet them. The resulting texture (Figure 37a) consists of three noticeable regions, which are highlighted with different color yarns. The grey region at the top and bottom of the fabric consists of all purl stitches. This causes both ends to curl backwards, making it easier to seam the fabric to the lamp shade frame.

Table 18: Lampshade Heuristic-Map

Objective Weight	Objective	modifier	modifier Weight
4	Grey Curled-Edges Region	Flip Stitch	1
3	Low-Opacity Blue Regions	Replace with Increase Decrease Pair	1
2	Maximize Horizontal Elasticity	Flip Stitch	2
		Replace with Cable Stitch	1
1	Cable Green Regions	Replace with Cable Stitch	2
		Replace with Inc Slip-Dec Pair	1

The green region consists of vertical stripes of cable stitches. The stripes mostly alternate in the lean direction, giving the appearance of ropes, though they are not as consistent as a designer might manually have selected. The light blue regions consists of centered, slipped, triple decreases and two yarnovers. This particular decrease looks similar to cable-stitches, but the added yarnovers allow more light to shine through. We can use the same approach, of mapping regions of texture to different objectives, to create other knit objects such as a crutch cosy (Figure 37b) or a colorwork pillow (Figure 37c).

10.5 CONCLUSION

KnitGIST demonstrates the flexibility of OPTIMUM as a meta-tool for implementing generative design tools. While OPTIMUM is intended to support the construction of clinical CAD tools, its support for orthogonal experts generalizes to a variety of domains. In this case, generative design for machine knitting is managed by applying domain expertise that is common among hand-knitters. Hand knitters recognize the relationship between different combination of stitches and texture properties. This knowledge is distilled into KnitGIST with objectives and modifiers, and can be recombined with heuristic maps to create a wide variety of textures.

11 CONCLUSIONS AND FUTURE WORK

This thesis presents a future where healthcare practitioners and people with disabilities can build and customize physical devices at the point-of-care, rather than using inflexible mass-manufactured devices. Transforming broader healthcare practice requires novel digital fabrication tools that enable diverse domain-experts (e.g., clinicians, disabled people, designers) to design, manufacture, and customize new devices. My research shows that, rather than adapting to generalized design tools, clinicians and people with disabilities need domain-specific tools that automate a variety of design tasks (e.g., adjusting parameters, adapting to machine requirements). To make building such a wide variety of tools feasible, my work presents domain-agnostic meta-tools (e.g., Chapter 7, Chapter 8) that enable teams of programmers and domain-experts to quickly build domain-specific design tools. These meta-tools make it easier to build innovative design tools for healthcare, accessibility, and other digital fabrication domains.

The systems I present in this thesis are primarily targeted at supporting rapid prototyping by clinicians and disabled people. However, the space of medical making raises many new challenges for digital fabrication research. In particular, my future work will focus on three areas. First, we must address clinician's concerns about the safety of medical making. Once they can create designs, they need systems that can verify that those designs will meet patients' needs and are safe to use. Second, these systems largely focus on supporting a single designer in creating a single device. But often, customization and design are collaborative processes that, in the context of healthcare and access, must include disabled collaborators. Making fabrication accessible, through a focus on collaboration, opens promising avenues for research. Finally, beyond medical and accessibility contexts, this work reveals the numerous challenges of translating a prototype into a manufacturable product, be it a splint, face shield, or knitted garment. Translating digital fabrication technologies from prototyping to manufacturing tools opens opportunities to explore novel localized and sustainable fabrication practices.

11.1 VERIFICATION AND SAFETY FOR PATIENT SPECIFIC MEDICAL MAKING

Device safety and clinician trust were consistent themes in my studies of medical making; it was the principle concern of nearly every clinician I engaged with. Some of these concerns arose from a lack of familiarity with digital fabrication technologies. How can a clinician who is trained in assessing the status of a human body evaluate the mechanical properties of a 3D printed device? It is the role of technical collaborators (e.g., engineers) and design tools to help clinicians understand risks. However, customizing designs to a specific person raises another question that engineers and design tools cannot readily resolve; which changes to a standardized design pose unknown and/or significant risks to the patient? Adding 5mm to the width of a splint wing appears insignificant, but could cause a dangerous abrasion on the patient's skin or increase the splint's risk of shattering. Changing the 3D printing material of an N95 respirator design appeared benign to many medical makers, including clinicians, but had substantial effects on the size of particles (e.g., viruses) filtered by the mask. In some cases, such as the N95 masks, the original designers could anticipate the introduced risks, however, in other cases (e.g., splints), the risks are patient-specific and need to be assessed by the clinician and patient by testing multiple devices.

While assessing all potential risks of any configuration of a customized device is likely infeasible, estimation of risk can be aided by growing data sets of medical device designs. The NIH 3D print exchange now offers a unique set of over 1000 designs, that have been clinically reviewed by experts and offer substantial documentation (e.g., 3D models, manufacturing instructions, testing data, usage instructions). By viewing repositories of designs as data sets, rather than assessing designs individually, we are presented with opportunities to apply new technologies (e.g., machine learning) to risk assessment problems. However, adapting these repositories to those methods will require modifications to the entire tool chain: creating fully documented modifiable designs (e.g., Chapter 7), classifications of relevant modifications and evaluation criteria by domain experts (e.g., Chapter 8), and improvements to end-user machine learning that enables domain experts to explore models that evaluate and understand device risks.

I expect new systems for evaluating risk will go beyond medical making to traditional medical device manufacturing and regulation. In the United States, FDA approval processes are used to assess the risks of novel medical devices. Often, this requires clinical studies of medical devices. However, the majority of devices are approved without providing "*reasonable assurance of safety and effectiveness for the device*" (e.g., clinical trials), because of their similarity to "*devices that have existing or reasonably foreseeable characteristics of commercially distributed devices*" [48] under FDA 510k exceptions. Essentially, a commercial medical device can be released to the public if the manufacturer argues that it is not substantially different from an already approved device. The 510k exemption is critical to the regulatory future of patient-specific digital fabrication, since this is what allows modifications of devices in medical settings. That said, the exemption has also lead to the approval of multiple devices that were eventually removed from the market after causing significant harm, including death, to patients [173]. For example, cobalt hip replacements were approved without clinical studies due to their similarity to ceramic hip replacements of the same shape. However, studies have shown that the cobalt, when damaged from overuse or a fall, can get into the bloodstream and cause a variety significant health concerns (e.g., deafness, hypothyroidism, tremor, dementia, heart failure). Using cobalt appeared to FDA regulators as a benign change, but cobalt had never been implanted near muscle tissue in any other FDA approved device. Systems that help medical device designers and regulators examine the differences between novel designs and existing designs have the potential to formalize 510k exceptions, assuring the future of medical making, while limiting and elucidating the significant risks introduced by medical devices.

RESEARCH QUESTIONS:

- How do clinicians understand device risks when devices are changed and modified to fit a patient? How is this affected by outside measures (e.g., FDA Approval) and patient's co-morbidities?
- How can design tools enable designers, including non-technical experts, to evaluate and measure potential risks from design modifications?
- Can large data sets of medical device designs be used to evaluate a novel device's similarity to existing designs and risk profile?

11.2 ACCESSIBLE AND COLLABORATIVE FABRICATION

So much of the experience of disability is a mismatch between the body and the built environment. This is why digital fabrication has such significant potential as a tool to increase accessibility, by reducing the complexity of modifying the physical world to meet individual needs. However, access cannot be assured until all stakeholders are able to use digital fabrication tools. My focus on medical making expands the tool set of clinicians, one common provider of assistive technology. Still, many other stakeholders are needed, particularly people with disabilities themselves. Prior work has explored increasing the accessibility of digital fabrication tools primarily for people who are blind or have low-vision [228]. Future work must build on these efforts to increase digital fabrication access for all.

The core challenge of increasing access to digital fabrication technologies is making the full complexity of design options available through systems that accommodate different levels of abilities. I argue that tackling individual impairments—e.g., limited mobility making 3D modeling difficult, limited vision impairing access to 3D models—is an ineffective way of addressing wider access to digital fabrication. The range of human abilities and disability is wider than we can address with adapted tools. However, collaborative teams are often able to overcome an individual's impairments.

Instead of focusing on making digital fabrication tools accessible to individual impairments, we can approach access through an interdependence lens [24]. Disabled people surround themselves with mutual aid structures that enable them to solve access challenges. Digital fabrication can be one more tool available to these communities to solve those challenges. However, it requires that CAD tools be reconstructed to accommodate collaboration by groups with diverse and mixed abilities. Consider a design team working with a blind person. A sighted 3D modeler can do most of the work of creating models from the blind person's description, however the BLV person still needs a tool to evaluate the model (e.g., [228]) and offer critique. Further, by focusing on collaborative design tasks, this research will have to directly tackle access conflicts when multiple collaborators have disabilities. Rather than creating an assistive solution for one person with one type of impairment, we must create these tools to accommodate an ecosystem of diverse disabilities, rather than individual use cases.

Exploring digital fabrication as a collaborative tool for provisioning access offers a unique test bed for a new focus of accessibility research that facilitates communities actualizing access in-situ, rather than addressing individual impairments. Viewing accessible digital fabrication as a challenge of collaboration raises a variety of questions for future work:

RESEARCH QUESTIONS:

- How can collaborative CAD tools be built to assure equitable agency in the design process among collaborators? Particularly when power dynamics may affect equity such as: patient and clinician, student and instructor, care giver and recipient?
- Can design tools and artifacts be used as a medium to communicate in a collaborative design process when collaborators use different communication methods (e.g., speechreading and sign language, AAC devices) or have different types of expertise?
- How can designers more effectively collaborate with AI agents (e.g., generative design tools) as an alternative modality for design?

11.3 SUSTAINABLE AND LOCALIZED MANUFACTURING

Digital fabrication, particularly when focused on 3D printing technologies, has largely been focused on supporting rapid prototyping—getting to the right design. Advanced manufacturing research has focused on the equally critical task of enabling production more complex designs. However, supporting the production of the right designs remains a critical gap [99]. Particularly in the context of medical making, this gap makes it difficult to transfer medical device prototypes to medical practice. The occupational therapists in Chapter 3 saw this as the primary barrier to establishing trust in the safety of their devices. When producing PPE for the pandemic, makers in a variety of contexts struggled to reliably manufacture 3D models on their 3D printers. In the space of knitwear manufacturing, the gap between design tools and production tools is even wider and accounts for significant waste in the textile industry [270]. The space between prototyping and production is a promising and broad area of future work.

The COVID-19 pandemic has caused numerous failures in global supply chains. Two years after the start of the pandemic, many products remain difficult to procure or delayed. The effects of these failures are most severely felt by low-resource communities that have less flexibility in adapting to changing circumstance. Just as digital fabrication stepped in to replace traditional PPE supply chains, it may play a role in supporting diverse communities in the future. In particular, digital fabrication technologies can empower communities and individuals to produce for themselves, locally. However, this will require supporting new groups of users that have, like disabled people, often been missed when developing CAD tools for engineers.

Beyond empowerment, localized supply chains can increase sustainability. At a minimum, they remove the significant costs of transportation and over-production. Multiple knitwear startups are already jumping into this space, using knitting machines as a way to locally produce garments on demand. However, the lack of prototyping-to-production knitting design tools imposes a significant burden on the knit designers leading this new industry. I believe this makes knit-programming a promising test bed for new design tools that support both prototyping and production.

Branching from my focus on health and accessibility, this research in sustainable and localized manufacturing raises exciting questions for future work.

RESEARCH QUESTIONS:

- How can advances in Knit-Programming enable end-user customized, localized, and sustainable manufacturing of textiles and garments?
- How can design tools support rapid prototyping design practices that lead to ready-to-manufacture designs?
- What is the role of digital fabrication technologies in crisis supply chains?

11.4 EXPANDING THE FLEXIBILITY OF COMPUTING

Computing is a highly flexible material for problem solving and design. We can rapidly build, test, modify, and deploy solutions at massive scales. Once out in the world, those solutions are not static. They can be further modified to solve highly individualized challenges. This makes it a critical tool in solving societal challenges of health and accessibility, which are both influenced by social infrastructure and individual circumstances. My studies and collaborations in a variety of healthcare settings have shown the challenges clinicians and people with disabilities face when solving physical problems (e.g., assistive device design, medical device supply chains) without the flexibility of computing solutions.

Digital fabrication tools bridge the space between digital and physical design problems, carrying the benefits of both, but also their inherit challenges. Digital mediums are highly reusable—shared globally at the click of mouse. This also makes it difficult to ensure changes in designs do not get out of hand and put users at risk. Similarly, the physical products based on those designs have to be made in different physical contexts (e.g., factories, hospitals, home makerspaces) that each carry unique physical limitations and capabilities. Meta-tools like PARTs apply strategies from software engineering to the design of physical objects, making their transfer to new contexts and users easier and more reliable. Alternatively, OPTIMUM makes use of a key advantage of computing, iteration speed, to resolve tasks that are often left to designers. It does this in a variety of domains, by including domain experts in the optimization design process. The work of OPTIMUM shows that programming expertise alone is insufficient for building tools that modify the physical world. We need the expertise of clinicians to describe clinical decision making processes (e.g., cataract lens selection, splint design). We need people with disabilities to describe what access solutions work best for them (e.g., Maptimizer). We need knit-designers to define the best ways to turn a bundle of yarn into functional and attractive fabric objects. Rather than educating different types of users about how to program, or simplifying computing tasks to make them generally accessible to novices, my work takes the flexibility of computing and adapts it to the established practices of different communities.

BIBLIOGRAPHY

[1] In: *ACM Trans. Graph.* 40.6 (2021). ISSN: 0730-0301.

[2] *2019 SOLIDWORKS Help - Smart Fasteners Overview*. [Online; accessed 7. Dec. 2020]. 2020. URL: https://help.solidworks.com/2019/english/SolidWorks/sldworks/c_Smart_Fasteners_Overview.htm.

[3] *3D Printing Medical Devices at the Point of Care: Discussion Paper*. Tech. rep. U.S. Food and Drug Administration, Oct. 2021. URL: https://www.fda.gov/medical-devices/3d-printing-medical-devices/3d-printing-medical-devices-point-care-discussion-paper.

[4] *A Comparison of Cooling Schedules for Simulated Annealing (Artificial Intelligence)*. [Online; accessed 8. Feb. 2022]. 2022. URL: http://what-when-how.com/artificial-intelligence/a-comparison-of-cooling-schedules-for-simulated-annealing-artificial-intelligence.

[5] Muhammad Abdullah, Martin Taraz, Yannis Kommana, Shohei Katakura, Robert Kovacs, Jotaro Shigeyama, Thijs Roumen, and Patrick Baudisch. "FastForce: Real-Time Reinforcement of Laser-Cut Structures." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. URL: https://doi.org/10.1145/3411764.3445466.

[6] Mathieu Acher, Benoit Baudry, Olivier Barais, and Jean-Marc Jezequel. "Customization and 3D Printing: A Challenging Playground for Software Product Lines." In: *Proceedings of the 18th International Software Product Line Conference - Volume 1*. SPLC '14. ACM, 2014, pp. 142–146. ISBN: 978-1-4503-2740-4. DOI: 10.1145/2648511.2648526. URL: http://doi.acm.org/10.1145/2648511.2648526.

[7] L. Aflatoony and S. J. Lee. "CODEA: A Framework for Co-Designing Assistive Technologies with Occupational Therapists, Industrial Designers, and End-Users with Mobility Impairments." In: *Proceedings of the Design Society: DESIGN Conference* 1 (2020), pp. 1843–1852. ISSN: 2633-7762. DOI: 10.1017/dsd.2020.332.

[8] Leila Aflatoony and Su Jin (Susan) Lee. "AT Makers: A Multidisciplinary Approach to Co-Designing Assistive Technologies by Co-Optimizing Expert Knowledge." In: *Proceedings of the 16th Participatory Design Conference 2020 - Participation(s) Otherwise - Volume 2*. PDC '20. Association for Computing Machinery, 2020, 128–132. ISBN: 9781450376068. DOI: 10.1145/3384772.3385158. URL: https://doi.org/10.1145/3384772.3385158.

[9] Asterios Agkathidis. "Generative Design Methods." In: vol. 2. 2015.

[10] Maneesh Agrawala and Chris Stolte. "Rendering Effective Route Maps: Improving Usability through Generalization." In: *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '01. New York, NY, USA: Association for Computing Machinery, 2001, 241–249. ISBN: 158113374X. DOI: 10.1145/383259.383286. URL: https://doi.org/10.1145/383259.383286.

[11] Lea Albaugh, Scott Hudson, and Lining Yao. "Digital Fabrication of Soft Actuated Objects by Machine Knitting." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2019. DOI: 10.1145/3290605.3300414.

[12] Lea Albaugh, James McCann, Scott E. Hudson, and Lining Yao. "Engineering Multifunctional Spacer Fabrics Through Machine Knitting." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445564. URL: https://doi.org/10.1145/3411764.3445564.

[13] Meryl Alper. "Making space in the makerspace: Building a mixed-ability maker culture." In: *Proceedings of the Interaction Design and Children (IDC-13), New York, NY, USA* (2013), pp. 24–27.

[14] Fraser Anderson, Tovi Grossman, and George Fitzmaurice. "Trigger-Action-Circuits: Leveraging Generative Design to Enable Novices to Design and Build Circuitry." In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. UIST '17. Québec City, QC, Canada: Association for Computing Machinery, 2017, 331–342. ISBN: 9781450349819. DOI: 10.1145/3126594.3126637. URL: https://doi.org/10.1145/3126594.3126637.

[15] M. D. Arthur M. Dannenberg. "A Simple Face mask for use by Contagious Disease Attendants." In: *JAMA* 70.14 (1918), p. 990. ISSN: 0002-9955. DOI: 10.1001/jama.1918.26010140001011.

[16] Moritz Bächer, Bernd Bickel, Doug L. James, and Hanspeter Pfister. "Fabricating Articulated Characters from Skinned Meshes." In: *ACM Trans. Graph.* 31.4 (2012). ISSN: 0730-0301. DOI: 10.1145/2185520.2185543. URL: https://doi.org/10.1145/2185520.2185543.

[17] Moritz Bächer, Stelian Coros, and Bernhard Thomaszewski. "LinkEdit: Interactive Linkage Editing Using Symbolic Kinematics." In: *ACM Trans. Graph.* 34.4 (2015). ISSN: 0730-0301. DOI: 10.1145/2766985. URL: https://doi.org/10.1145/2766985.

[18] Moritz Bächer, Emily Whiting, Bernd Bickel, and Olga Sorkine-Hornung. "Spin-It: Optimizing Moment of Inertia for Spinnable Objects." In: *ACM Trans. Graph.* 33.4 (2014). ISSN: 0730-0301. DOI: 10.1145/2601097.2601157. URL: https://doi.org/10.1145/2601097.2601157.

[19] David Bak. "Rapid prototyping or rapid production? 3D printing processes move industry towards the latter." In: *Assembly Automation* 23.4 (2003), pp. 340–345.

[20] Sasha Barab and Kurt Squire. "Design-Based Research: Putting a Stake in the Ground." In: *Journal of the Learning Sciences* 13.1 (2004), pp. 1–14. ISSN: 1050-8406. DOI: 10.1207/s15327809jls1301_1. URL: https://doi.org/10.1207/s15327809jls1301_1.

[21] Jay M. Baruch. "Doctors as Makers." In: *Academic Medicine* 92.1 (2017). ISSN: 1040-2446. URL: https://journals.lww.com/academicmedicine/Fulltext/2017/01000/Doctors_as_Makers.17.aspx.

[22] Jonathan Bean and Daniela Rosner. "Making: Movement or Brand?" In: *interactions* 21.1 (2014), pp. 26–27. ISSN: 1072-5520. DOI: 10.1145/2541669. URL: http://doi.acm.org/10.1145/2541669.

[23] M.P. Bendsoe and O. Sigmund. *Topology Optimization: Theory, Methods, and Applications*. Engineering online library. Springer Berlin Heidelberg, 2003. ISBN: 9783540429920. URL: https://books.google.com/books?id=NGmtmMhVe2sC.

[24] Cynthia L. Bennett, Erin Brady, and Stacy M. Branham. "Interdependence as a Frame for Assistive Technology Research and Design." In: *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '18. Association for Computing Machinery, 2018, 161–173. ISBN: 9781450356503. DOI: 10.1145/3234695.3236348. URL: https://doi.org/10.1145/3234695.3236348.

[25] Cynthia L. Bennett, Keting Cen, Katherine M. Steele, and Daniela K. Rosner. "An Intimate Laboratory?: Prostheses As a Tool for Experimenting with Identity and Normalcy." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. ACM, 2016, pp. 1745–1756. ISBN: 978-1-4503-3362-7. DOI: 10.1145/2858036.2858564. URL: http://doi.acm.org/10.1145/2858036.2858564.

[26] Cynthia L. Bennett, Burren Peil, and Daniela K. Rosner. "Biographical Prototypes: Reimagining Recognition and Disability in Design." In: *Proceedings of the 2019 on Designing Interactive Systems Conference*. DIS '19. Association for Computing Machinery, 2019, 35–47. ISBN: 9781450358507. DOI: 10.1145/3322276.3322376. URL: https://doi.org/10.1145/3322276.3322376.

[27] Cynthia L. Bennett, Abigale Stangl, Alexa F. Siu, and Joshua A. Miele. "Making Nonvisually: Lessons from the Field." In: *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '19. Association for Computing Machinery, 2019, 279–285. ISBN: 9781450366762. DOI: 10.1145/3308561.3355619. URL: https://doi.org/10.1145/3308561.3355619.

[28] Gaurav Bharaj, David I. W. Levin, James Tompkin, Yun Fei, Hanspeter Pfister, Wojciech Matusik, and Changxi Zheng. "Computational Design of Metallophone Contact Sounds." In: *ACM Trans. Graph.* 34.6 (2015). ISSN: 0730-0301. DOI: 10.1145/2816795.2818108. URL: https://doi.org/10.1145/2816795.2818108.

[29] Gina Bianchini. *Home*. en-US. 2017. URL: https://www.mightynetworks.com/new.

[30] Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. "Design and Fabrication of Materials with Desired Deformation Behavior." In: *ACM Trans. Graph.* 29.4 (2010). ISSN: 0730-0301. DOI: 10.1145/1778765.1778800. URL: https://doi.org/10.1145/1778765.1778800.

[31] Christian Blum and Andrea Roli. "Metaheuristics in Combinatorial Optimization: Overview and Conceptual Comparison." In: *ACM Comput. Surv.* 35.3 (2003), 268–308. ISSN: 0360-0300. DOI: 10.1145/937503.937505. URL: https://doi.org/10.1145/937503.937505.

[32] Christian Blum and Andrea Roli. "Metaheuristics in Combinatorial Optimization: Overview and Conceptual Comparison." In: *ACM Comput. Surv.* 35.3 (2003), 268–308. ISSN: 0360-0300. DOI: 10.1145/937503.937505. URL: https://doi.org/10.1145/937503.937505.

[33] JC Briar. *Stitch Maps*. 2013. URL: https://stitch-maps.com/.

[34] Craig Brown and Amy Hurst. "VizTouch: Automatically Generated Tactile Visualizations of Coordinate Spaces." In: *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*. TEI '12. ACM, 2012, pp. 131–138. ISBN: 978-1-4503-1174-8. DOI: 10.1145/2148131.2148160. URL: http://doi.acm.org/10.1145/2148131.2148160.

[35] Christopher Buccafusco. *Disability and Design*. [Online; accessed 6. Nov. 2020]. 2020. DOI: 10.2139/ssrn.3497902.

[36] Erin Buehler, Stacy Branham, Abdullah Ali, Jeremy J. Chang, Megan Kelly Hofmann, Amy Hurst, and Shaun K. Kane. "Sharing is Caring: Assistive Technology Designs on Thingiverse." In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. ACM, 2015, pp. 525–534. ISBN: 978-1-4503-3145-6. DOI: 10.1145/2702123.2702525. URL: http://doi.acm.org/10.1145/2702123.2702525.

[37] Erin Buehler, Niara Comrie, Megan Hofmann, Samantha McDonald, and Amy Hurst. "Investigating the Implications of 3D Printing in Special Education." In: *ACM Trans. Access. Comput.* 8.3 (2016), 11:1–11:28. ISSN: 1936-7228. DOI: 10.1145/2870640. URL: http://doi.acm.org/10.1145/2870640.

[38] Erin Buehler, Amy Hurst, and Megan Hofmann. "Coming to Grips: 3D Printing for Accessibility." In: *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility*. ASSETS '14. ACM, 2014, pp. 291–292. ISBN: 978-1-4503-2720-6. DOI: 10.1145/2661334.2661345. URL: http://doi.acm.org/10.1145/2661334.2661345.

[39] Matthew B. Burn, Anderson Ta, and Gloria R. Gogola. "Three-Dimensional Printing of Prosthetic Hands for Children." In: *Journal of Hand Surgery* 41.5 (2016), e103–e109. ISSN: 0363-5023. DOI: 10.1016/j.jhsa.2016.02.008.

[40] Edward Burnell, Nicole B. Damen, and Warren Hoburg. "GPkit: A Human-Centered Approach to Convex Optimization in Engineering Design." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2020, 1–13. ISBN: 9781450367080. URL: https://doi.org/10.1145/3313831.3376412.

[41] Bill Buxton. *Sketching User Experiences: Getting the Design Right and the Right Design*. 1st Edition. 2007. ISBN: 978-0-08-055290-3.

[42] *CHI '21: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966.

[43] Scott Carter, Amy Hurst, Jennifer Mankoff, and Jack Li. "Dynamically Adapting GUIs to Diverse Input Devices." In: *Proceedings of the 8th International ACM SIGACCESS Conference on Computers and Accessibility*. Assets '06. Association for Computing Machinery, 2006, 63–70. ISBN: 1595932909. DOI: 10.1145/1168987.1169000.

[44] Weikai Chen, Yuexin Ma, Sylvain Lefebvre, Shiqing Xin, Jonàs Martínez, and wenping wang. "Fabricable Tile Decors." In: *ACM Trans. Graph.* 36.6 (2017). ISSN: 0730-0301. DOI: 10.1145/3130800.3130817. URL: https://doi.org/10.1145/3130800.3130817.

[45] Xiang 'Anthony' Chen, Jeeeun Kim, Jennifer Mankoff, Tovi Grossman, Stelian Coros, and Scott E. Hudson. "Reprise: A Design Tool for Specifying, Generating, and Customizing 3D Printable Adaptations on Everyday Objects." In: *Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology*. Tokyo, Japan: ACM, 2016. ISBN: 978-1-4503-4189-9. DOI: 10.1145/2984511.2984512. URL: http://dx.doi.org/10.1145/2984511.2984512.

[46] Xiang "Anthony" Chen, Ye Tao, Guanyun Wang, Runchang Kang, Tovi Grossman, Stelian Coros, and Scott E. Hudson. "Forte: User-Driven Generative Design." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Association for Computing Machinery, 2018. ISBN: 9781450356206. DOI: 10.1145/3173574.3174070. URL: https://doi.org/10.1145/3173574.3174070.

[47] Gabriel Cirio, Jorge Lopez-Moreno, David Miraut, and Miguel A. Otaduy. "Yarn-level Simulation of Woven Cloth." In: *ACM Trans. Graph.* 33.6 (2014), 207:1–207:11. ISSN: 0730-0301. DOI: 10.1145/2661229.2661279. URL: http://doi.acm.org/10.1145/2661229.2661279.

[48] *Class I and Class II Device Exemptions*. [Online; accessed 18. Apr. 2022]. Feb. 2022. URL: https://www.fda.gov/medical-devices/classify-your-medical-device/class-i-and-class-ii-device-exemptions.

[49] Christina Cogdell. *Toward a Living Architecture?: Complexism and Biology in Generative Design*. U of Minnesota Press, 2018. ISBN: 9781452958071.

[50] Stelian Coros, Bernhard Thomaszewski, Gioacchino Noris, Shinjiro Sueda, Moira Forberg, Robert W. Sumner, Wojciech Matusik, and Bernd Bickel. "Computational Design of Mechanical Characters." In: *ACM Trans. Graph.* 32.4 (2013). ISSN: 0730-0301. DOI: 10.1145/2461912.2461953. URL: https://doi.org/10.1145/2461912.2461953.

[51] John W. Creswell and Dana L. Miller. "Determining Validity in Qualitative Inquiry." In: *Theory Into Practice* 39.3 (2000). Publisher: Routledge _eprint: https://doi.org/10.1207/s15430421tip3903_2, pp. 124–130. ISSN: 0040-5841. DOI: 10.1207/s15430421tip3903_2. URL: https://doi.org/10.1207/s15430421tip3903_2.

[52] Sebastian Cucerca, Piotr Didyk, Hans-Peter Seidel, and Vahid Babaei. "Computational Image Marking on Metals via Laser Induced Heating." In: *ACM Trans. Graph.* 39.4 (2020). ISSN: 0730-0301. DOI: 10.1145/3386569.3392423. URL: https://doi.org/10.1145/3386569.3392423.

[53] Maitraye Das, Katya Borgos-Rodriguez, and Anne Marie Piper. "Weaving by Touch: A Case Analysis of Accessible Making." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20. Association for Computing Machinery, 2020, 1–15. ISBN: 9781450367080. DOI: 10.1145/3313831.3376477. URL: https://doi.org/10.1145/3313831.3376477.

[54] Lieven De Couvreur and Richard Goossens. "Design for (every) one: co-creation as a bridge between universal design and rehabilitation engineering." In: *CoDesign* 7.2 (2011), pp. 107–121.

[55] Jack B. Dennis. "Modularity." In: *Software Engineering, An Advanced Course, Reprint of the First Edition [February 21 - March 3, 1972]*. Berlin, Heidelberg: Springer-Verlag, 1975, 128–182. ISBN: 3540071687.

[56] Ruta Desai, James McCann, and Stelian Coros. "Assembly-Aware Design of Printable Electromechanical Devices." In: *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST '18. Association for Computing Machinery, 2018, 457–472. ISBN: 9781450359481. DOI: 10.1145/3242587.3242655. URL: https://doi.org/10.1145/3242587.3242655.

[57] Yue Dong, Jiaping Wang, Fabio Pellacini, Xin Tong, and Baining Guo. "Fabricating Spatially-Varying Subsurface Scattering." In: *ACM SIGGRAPH 2010 Papers*. SIGGRAPH '10. Association for Computing Machinery, 2010. ISBN: 9781450302104. DOI: 10.1145/1833349.1778799. URL: https://doi.org/10.1145/1833349.1778799.

[58] M. Dorigo, M. Birattari, and T. Stutzle. "Ant colony optimization." In: *IEEE Computational Intelligence Magazine* 1.4 (2006), pp. 28–39. DOI: 10.1109/MCI.2006.329691.

[59] Steven P. Dow, Kate Heddleston, and Scott R. Klemmer. "The Efficacy of Prototyping Under Time Constraints." In: *Proceedings of the Seventh ACM Conference on Creativity and Cognition*. C&C '09. ACM, 2009, pp. 165–174. ISBN: 978-1-60558-865-0. DOI: 10.1145/1640233.1640260. URL: http://doi.acm.org/10.1145/1640233.1640260.

[60] *DtM-v3.1 Face Shield PPE, 3D printable headband No logo*. 2020. URL: https://3dprint.nih.gov/discover/3dpx-013359.

[61] Jérémie Dumas, Jean Hergel, and Sylvain Lefebvre. "Bridging the Gap: Automated Steady Scaffoldings for 3D Printing." In: *ACM Trans. Graph.* 33.4 (2014). ISSN: 0730-0301. DOI: 10.1145/2601097.2601153. URL: https://doi.org/10.1145/2601097.2601153.

[62] David J. Earl and Michael W. Deem. "Parallel tempering: Theory, applications, and new perspectives." In: *Phys. Chem. Chem. Phys.* 7.23 (2005), pp. 3910–3916. DOI: 10.1039/B509983H.

[63] Rafik Elmansy. "Designing the 3D-Printed Prosthetic Hand." In: *Design Management Review* 26.1 (2015), pp. 24–31.

[64] Adam EM Eltorai, Eric Nguyen, and Alan H Daniels. "Three-dimensional printing in orthopedic surgery." In: *Orthopedics* 38.11 (2015), pp. 684–687.

[65] Michael T. M. Emmerich and André H. Deutz. "A tutorial on multiobjective optimization: fundamentals and evolutionary methods." In: *Nat. Comput.* 17.3 (2018), pp. 585–609. ISSN: 1572-9796. DOI: 10.1007/s11047-018-9685-y.

[66] *Enabling The Future*. en-US. 2014. URL: https://enablingthefuture.org/.

[67] M Angeles Espinosa, Simon Ungar, Esperanza Ochaíta, Mark Blades, and Christopher Spencer. "Comparing methods for introducing blind and visually impaired people to unfamiliar urban environments." In: *Journal of environmental psychology* 18.3 (1998), pp. 277–287.

[68] Frederika Farley. "Improvising Equipment." In: *The American Journal of Nursing* (1938), 42s–43s.

[69] Blake Farmer. *FDA Tightening Regulatory Requirements For Some Medical Devices*. 2019. URL: https://www.wbur.org/npr/689739642/fda-tightening-regulatory-requirements-for-some-medical-devices.

[70] Gary Alan Fine. "The Sociology of the Local: Action and its Publics*." In: *Sociological Theory* 28.4 (2010), pp. 355–376. ISSN: 0735-2751. DOI: 10.1111/j.1467-9558.2010.01380.x.

[71] Joseph Flynt. *3D Printing Hospital Network to be Developed by VA*. en-US. 2017. URL: https://3dinsider.com/3d-printing-hospital-va/.

[72] James Fogarty and Scott E. Hudson. "GADGET: A Toolkit for Optimization-Based Approaches to Interface and Display Generation." In: UIST '03 (2003), 125–134. DOI: 10.1145/964696.964710. URL: https://doi.org/10.1145/964696.964710.

[73] Sarah Fox, Rachel Rose Ulgado, and Daniela Rosner. "Hacking Culture, Not Devices: Access and Recognition in Feminist Hackerspaces." In: *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. CSCW '15. ACM, 2015, pp. 56–68. ISBN: 978-1-4503-2922-4. DOI: 10.1145/2675133.2675223. URL: http://doi.acm.org/10.1145/2675133.2675223.

[74] Krzysztof Z. Gajos, Daniel S. Weld, and Jacob O. Wobbrock. "Automatically generating personalized user interfaces with Supple." In: *Artif. Intell.* 174.12 (2010), pp. 910–950. ISSN: 0004-3702. DOI: 10.1016/j.artint.2010.05.005.

[75] Krzysztof Gajos and Daniel S. Weld. "SUPPLE: Automatically Generating User Interfaces." In: *Proceedings of the 9th International Conference on Intelligent User Interfaces*. IUI '04. Association for Computing Machinery, 2004, pp. 93–100. ISBN: 1581138156. DOI: 10.1145/964442.964461. URL: https://doi.org/10.1145/964442.964461.

[76] Akash Garg, Alec Jacobson, and Eitan Grinspun. "Computational Design of Reconfigurables." In: 35.4 (2016). ISSN: 0730-0301. DOI: 10.1145/2897824.2925900. URL: https://doi.org/10.1145/2897824.2925900.

[77] Emilie Giles, Janet van der Linden, and Marian Petre. "Weaving Lighthouses and Stitching Stories: Blind and Visually Impaired People Designing E-textiles." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM. 2018, p. 470.

[78] Stéphanie Giraud, Anke M Brock, Marc J-M Macé, and Christophe Jouffrais. "Map learning with a 3D printed interactive small-scale model: Improvement of space and text memorization in visually impaired students." In: *Frontiers in psychology* 8 (2017), p. 930.

[79] Mary Ellen Smith Glasgow, Alison Colbert, John Viator, and Stephen Cavanagh. "The Nurse-Engineer: A New Role to Improve Nurse Technology Interface and Patient Care Device Innovations." In: *Journal of Nursing Scholarship* 50.6 (2018), pp. 601–611. DOI: 10.1111/jnu.12431. eprint: https://sigmapubs.onlinelibrary.wiley.com/doi/pdf/10.1111/jnu.12431. URL: https://sigmapubs.onlinelibrary.wiley.com/doi/abs/10.1111/jnu.12431.

[80] Ephraim P. Glinert and Bryant W. York. "Computers and people with disabilities." In: *Commun. ACM* 35.5 (1992), pp. 32–35. ISSN: 0001-0782. DOI: 10.1145/129875.129876.

[81] Fred Glover and Manuel Laguna. "Tabu Search." In: *Handbook of Combinatorial Optimization: Volume1–3*. Ed. by Ding-Zhu Du and Panos M. Pardalos. Boston, MA: Springer US, 1998, pp. 2093–2229. ISBN: 978-1-4613-0303-9. DOI: 10.1007/978-1-4613-0303-9_33. URL: https://doi.org/10.1007/978-1-4613-0303-9_33.

[82] Sarah Anderson Goerhke. "Made in Gaza: Glia Project Offers Open Source, Low-Cost, Locally 3D Printed Medical Supplies." en-US. In: *3DPrint.com | The Voice of 3D Printing / Additive Manufacturing* (2015). URL: https://3dprint.com/95097/project-glia-gaza-strip/.

[83] Daniel Golovin, Benjamin Solnik, Subhodeep Moitra, Greg Kochanski, John Karro, and D. Sculley. "Google Vizier: A Service for Black-Box Optimization." In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD '17. New York, NY, USA: Association for Computing Machinery, 2017, 1487–1495. ISBN: 9781450348874. DOI: 10.1145/3097983.3098043. URL: https://doi.org/10.1145/3097983.3098043.

[84] Jose Gomez-Marquez and Anna Young. *A History of Nurse Making and Stealth Innovation*. Tech. rep. 2016. URL: http://dx.doi.org/10.2139/ssrn.2778663.

[85] Jun Gong, Olivia Seow, Cedric Honnet, Jack Forman, and Stefanie Mueller. "MetaSense: Integrating Sensing Capabilities into Mechanical Metamaterial." In: *The 34th Annual ACM Symposium on User Interface Software and Technology*. UIST '21. New York, NY, USA: Association for Computing Machinery, 2021, 1063–1073. ISBN: 9781450386357. DOI: 10.1145/3472749.3474806. URL: https://doi.org/10.1145/3472749.3474806.

[86] Timo Götzelmann. "LucentMaps: 3D Printed Audiovisual Tactile Maps for Blind and Visually Impaired People." In: *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '16. Association for Computing Machinery, 2016, 81–90. ISBN: 9781450341240. DOI: 10.1145/2982142.2982163. URL: https://doi.org/10.1145/2982142.2982163.

[87] Timo Götzelmann and Aleksander Pavkovic. "Towards automatically generated tactile detail maps by 3D printers for blind persons." In: *International Conference on Computers for Handicapped Persons*. Cham: Springer International Publishing, 2014, pp. 1–7.

[88] Floraine Grabler, Maneesh Agrawala, Robert W. Sumner, and Mark Pauly. "Automatic Generation of Tourist Maps." In: *ACM Trans. Graph.* 27.3 (2008), 1–11. ISSN: 0730-0301. DOI: 10.1145/1360612.1360699. URL: https://doi.org/10.1145/1360612.1360699.

[89] *Grasshopper - New in Rhino 6*. [Online; accessed 18. Nov. 2020]. 2020. URL: https://www.rhino3d.com/6/new/grasshopper.

[90] David Guirguis et al. "Evolutionary Black-Box Topology Optimization: Challenges and Promises." In: *IEEE Transactions on Evolutionary Computation* 24.4 (2020), pp. 613–633. DOI: 10.1109/TEVC.2019.2954411.

[91] Anhong Guo, Jeeeun Kim, Xiang 'Anthony' Chen, Tom Yeh, Scott E. Hudson, Jennifer Mankoff, and Jeffrey P. Bigham. "Facade: Auto-Generating Tactile Interfaces to Appliances." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. Association for Computing Machinery, 2017, 5826–5838. ISBN: 9781450346559. DOI: 10.1145/3025453.3025845. URL: https://doi.org/10.1145/3025453.3025845.

[92] Richard Guy and Khai Truong. "CrossingGuard: Exploring Information Content in Navigation Aids for Visually Impaired Pedestrians." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2012, 405–414. ISBN: 9781450310154. URL: https://doi.org/10.1145/2207676.2207733.

[93] Sara B. Hansen. "Tech's love affair with Colorado expected to continue despite global pandemic." en-US. In: *The Denver Post* (2020). URL: https://www.denverpost.com/2020/09/06/colorado-tech-boom-despite-coronavirus/.

[94] Björn Hartmann, Scott Doorley, and Scott R Klemmer. "Hacking, mashing, gluing: Understanding opportunistic design." In: *IEEE Pervasive Computing* 7.3 (2008), pp. 46–54.

[95] Mark R Hatch. *The maker revolution: Building a future on creativity and innovation in an exponential world*. John Wiley & Sons, 2017.

[96] Peregrine Hawthorn and Daniel Ashbrook. "Cyborg Pride: Self-Design in e-NABLE." In: *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '17. Association for Computing Machinery, 2017, 422–426. ISBN: 9781450349260. DOI: 10.1145/3132525.3134780. URL: https://doi.org/10.1145/3132525.3134780.

[97] Liang He, Zijian Wan, Leah Findlater, and Jon E. Froehlich. "TacTILE: A Preliminary Toolchain for Creating Accessible Graphics with 3D-Printed Overlays and Auditory Annotations." In: *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '17. Association for Computing Machinery, 2017, 397–398. ISBN: 9781450349260. DOI: 10.1145/3132525.3134818. URL: https://doi.org/10.1145/3132525.3134818.

[98] Nicholas Herbert, David Simpson, William D. Spence, and William Ion. "A preliminary investigation into the development of 3-D printing of prosthetic sockets." English. In: *Journal of Rehabilitation Research and Development* 42.2 (2005), pp. 141–6. URL: http://search.proquest.com/docview/215289926?accountid=9902.

[99] Steve Hodges. "Democratizing the Production of Interactive Hardware." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. New York, NY, USA: Association for Computing Machinery, 2020, 5–6. ISBN: 9781450375146. URL: https://doi.org/10.1145/3379337.3422877.

[100] Kenneth J. Hoffer. "The Hoffer Q formula: A comparison of theoretic and regression formulas." In: *Journal of Cataract & Refractive Surgery* 19.6 (1993), pp. 700–712. ISSN: 0886-3350. DOI: https://doi.org/10.1016/S0886-3350(13)80338-0. URL: https://www.sciencedirect.com/science/article/pii/S0886335013803380.

[101] Megan Hofmann, Julie Burke, Jon Pearlman, Goeran Fiedler, Andrea Hess, Jon Schull, Scott E. Hudson, and Jennifer Mankoff. "Clinical and Maker Perspectives on the Design of Assistive Technology with Rapid Prototyping Technologies." In: *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '16. ACM, 2016, pp. 251–256. ISBN: 978-1-4503-4124-0. DOI: 10.1145/2982142.2982181. URL: http://doi.acm.org/10.1145/2982142.2982181.

[102] Megan Hofmann, Jeffrey Harris, Scott E. Hudson, and Jennifer Mankoff. "Helping Hands: Requirements for a Prototyping Methodology for Upper-limb Prosthetics Users." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. ACM, 2016, pp. 1769–1780. ISBN: 978-1-4503-3362-7. DOI: 10.1145/2858036.2858340. URL: http://doi.acm.org/10.1145/2858036.2858340.

[103] Megan Hofmann, Jeffrey Harris, Scott E Hudson, and Jennifer Mankoff. "Helping Hands: Requirements for a Prototyping Methodology for Upper-limb Prosthetics Users." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM. 2016, pp. 1769–1780.

[104] Megan Hofmann, Devva Kasnitz, Jennifer Mankoff, and Cynthia L Bennett. "Living Disability Theory: Reflections on Access, Research, and Design." In: *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '20. New York, NY, USA: Association for Computing Machinery, 2020. ISBN: 9781450371032. DOI: 10.1145/3373625.3416996. URL: https://doi.org/10.1145/3373625.3416996.

[105] Megan Hofmann, Udaya Lakshmi, Kelly Mack, Scott E Hudson, Rosa I. Arriaga, and Jennifer Mankoff. "The Right to Help and the Right Help: Fostering and Regulating Collective Action in a Medical Making Reaction to COVID-19." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445707. URL: https://doi.org/10.1145/3411764.3445707.

[106] Megan Hofmann, Jennifer Mankoff, and Scott E. Hudson. "KnitGIST: A Programming Synthesis Toolkit for Generating Functional Machine-Knitting Textures." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST '20. Association for Computing Machinery, 2020, 1234–1247. ISBN: 9781450375146. DOI: 10.1145/3379337.3415590. URL: https://doi.org/10.1145/3379337.3415590.

[107] Megan Hofmann, Kristin Williams, Toni Kaplan, Stephanie Valencia, Gabriella Hann, Scott E. Hudson, Jennifer Mankoff, and Patrick Carrington. ""Occupational Therapy is Making": Clinical Rapid Prototyping and Digital Fabrication." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. ACM, 2019, 314:1–314:13. ISBN: 978-1-4503-5970-2. DOI: 10.1145/3290605.3300544. URL: http://doi.acm.org/10.1145/3290605.3300544.

[108] Megan Hofmann et al. "Maptimizer: Using Optimization to Tailor Tactile Maps to Users Needs." In: *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI '22. New York, NY, USA: Association for Computing Machinery, 2022. ISBN: 978-1-4503-9157-3/22/04. DOI: 10.1145/3491102.3517436. URL: https://doi.org/10.1145/3491102.3517436.

[109] Leona Holloway, Kim Marriott, and Matthew Butler. "Accessible Maps for the Blind: Comparing 3D Printed Models with Tactile Graphics." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 2018, 1–13. ISBN: 9781450356206. URL: https://doi.org/10.1145/3173574.3173772.

[110] Leona Holloway, Kim Marriott, and Matthew Butler. "Accessible Maps for the Blind: Comparing 3D Printed Models with Tactile Graphics." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2018, 1–13. ISBN: 9781450356206. URL: https://doi.org/10.1145/3173574.3173772.

[111] Leona Holloway, Kim Marriott, Matthew Butler, and Samuel Reinders. "3D Printed Maps and Icons for Inclusion: Testing in the Wild by People Who Are Blind or Have Low Vision." In: *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '19. New York, NY, USA: Association for Computing Machinery, 2019, 183–195. ISBN: 9781450366762. DOI: 10.1145/3308561.3353790. URL: https://doi.org/10.1145/3308561.3353790.

[112] K. Holtzblatt and H. Beyer. *Contextual Design: Defining Customer-Centered Systems*. Interactive Technologies. Elsevier Science, 1997. ISBN: 9780080503042. URL: https://books.google.com/books?id=JxQaQgOONGIC.

[113] Christopher Michael Homan, Jon I. Schull, and Akshai Prabhu. "On the Genesis of an Assistive Technology Crowdsourcing Community." In: *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. CHI EA '17. Association for Computing Machinery, 2017, 1003–1011. ISBN: 9781450346566. DOI: 10.1145/3027063.3053356. URL: https://doi.org/10.1145/3027063.3053356.

[114] Jonathan Hook, Sanne Verbaan, Abigail Durrant, Patrick Olivier, and Peter Wright. "A Study of the Challenges Related to DIY Assistive Technology in the Context of Children with Disabilities." In: *Proceedings of the 2014 Conference on Designing Interactive Systems*. DIS '14. ACM, 2014, pp. 597–606. ISBN: 978-1-4503-2902-6. DOI: 10.1145/2598510.2598530. URL: http://doi.acm.org/10.1145/2598510.2598530.

[115] Nathaniel Hudson, Celena Alcock, and Parmit K. Chilana. "Understanding Newcomers to 3D Printing: Motivations, Workflows, and Barriers of Casual Makers." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. ACM, 2016, pp. 384–396. ISBN: 978-1-4503-3362-7. DOI: 10.1145/2858036.2858266. URL: http://doi.acm.org/10.1145/2858036.2858266.

[116] Amy Hurst and Shaun Kane. "Making "Making" Accessible." In: *Proceedings of the 12th International Conference on Interaction Design and Children*. IDC '13. ACM, 2013, pp. 635–638. ISBN: 978-1-4503-1918-8. DOI: 10.1145/2485760.2485883. URL: http://doi.acm.org/10.1145/2485760.2485883.

[117] Amy Hurst and Jasmine Tobias. "Empowering Individuals with Do-it-yourself Assistive Technology." In: *The Proceedings of the 13th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '11. ACM, 2011, pp. 11–18. ISBN: 978-1-4503-0920-2. DOI: 10.1145/2049536.2049541. URL: http://doi.acm.org/10.1145/2049536.2049541.

[118] Darrell Hurt, Nick Weber, and Meghan Coakley McCarthy. *About the NIH 3D Print Exchange | NIH 3D Print Exchange*. URL: https://3dprint.nih.gov/about.

[119] *Instructables - How to make anything*. en. 2018. URL: https://www.instructables.com/.

[120] *Inworks CU Denver | Anschutz - An Innovation Lab for Learners & Leaders*. en-US. 2020. URL: https://inworks.ucdenver.edu/w/.

[121] Alexandra Ion, David Lindlbauer, Philipp Herholz, Marc Alexa, and Patrick Baudisch. "Understanding Metamaterial Mechanisms." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2019, 1–14. ISBN: 9781450359702. URL: https://doi.org/10.1145/3290605.3300877.

[122] Liz Jackson. "We Are the Original Lifehackers." In: *N.Y. Times* (2018). ISSN: 0362-4331. URL: https://www.nytimes.com/2018/05/30/opinion/disability-design-lifehacks.html.

[123] Dan Jacobson et al. *Spatial cognition through tactile mapping*. New York City, NY, USA: Cambridge University Press, 1992.

[124] Caigui Jiang, Chengcheng Tang, Hans-Peter Seidel, and Peter Wonka. "Design and Volume Optimization of Space Structures." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073619. URL: https://doi.org/10.1145/3072959.3073619.

[125] Danny L. Jorgensen. *Participant Observation: A Methodology for Human Studies*. en. Google-Books-ID: wi3ir38PfL4C. SAGE, 1989. ISBN: 978-0-8039-2877-0.

[126] Samuli Kärkkäinen. *Tactile Maps Easily | Touch Mapper*. [Online; accessed 25. Aug. 2021]. 2018. URL: https://touch-mapper.org/en.

[127] Deepak M. Kalaskar. "Challenges and Future Advancement." en. In: *3D Printing in Medicine*. Google-Books-ID: WrcxDQAAQBAJ. Woodhead Publishing, 2017, pp. 107–108. ISBN: 978-0-08-100726-6.

[128] Jonathan M. Kaldor, Doug L. James, and Steve Marschner. "Simulating Knitted Cloth at the Yarn Level." In: *ACM SIGGRAPH 2008 Papers*. SIGGRAPH '08. ACM, 2008, 65:1–65:9. ISBN: 978-1-4503-0112-1. DOI: 10.1145/1399504.1360664. URL: http://doi.acm.org/10.1145/1399504.1360664.

[129] Slim Kammoun, Florian Dramas, Bernard Oriolaand, and Christophe Jouffrais. "Route selection algorithm for blind pedestrian." In: *ICCAS 2010*. New York City, NY, USA: IEEE, 2010, pp. 2223–2228.

[130] Shaun K. Kane, Meredith Ringel Morris, and Jacob O. Wobbrock. "Touchplates: Low-Cost Tactile Overlays for Visually Impaired Touch Screen Users." In: *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '13. Association for Computing Machinery, 2013. ISBN: 9781450324052. DOI: 10.1145/2513383.2513442. URL: https://doi.org/10.1145/2513383.2513442.

[131] Alexandre Kaspar, Liane Makatura, and Wojciech Matusik. "Knitting Skeletons: A Computer-Aided Design Tool for Shaping and Patterning of Knitted Garments." In: *Proceedings of 32nd Annual ACM Symposium on User Interface Software and Technology*. UIST '19. New Orleans, LA: ACM, 2019.

[132] Alexandre Kaspar, Kui Wu, Yiyue Luo, Liane Makatura, and Wojciech Matusik. "Knit Sketching: From Cut and Sew Patterns to Machine-Knit Garments." In: *ACM Trans. Graph.* 40.4 (2021). ISSN: 0730-0301. DOI: 10.1145/3450626.3459752. URL: https://doi.org/10.1145/3450626.3459752.

[133] Jeeeun Kim, Anhong Guo, Tom Yeh, Scott E. Hudson, and Jennifer Mankoff. "Understanding Uncertainty in Measurement and Accommodating Its Impact in 3D Modeling and Printing." In: *Proceedings of the 2017 Conference on Designing Interactive Systems*. DIS '17. Association for Computing Machinery, 2017, 1067–1078. ISBN: 9781450349222. DOI: 10.1145/3064663.3064690. URL: https://doi.org/10.1145/3064663.3064690.

[134] Jeeeun Kim, Abigale Stangl, and Tom Yeh. "Using LEGO to Model 3D Tactile Picture Books by Sighted Children for Blind Children." In: *Proceedings of the 2nd ACM Symposium on Spatial User Interaction*. SUI '14. Association for Computing Machinery, 2014, p. 146. ISBN: 9781450328203. DOI: 10.1145/2659766.2661211. URL: https://doi.org/10.1145/2659766.2661211.

[135] Scott Kirkpatrick, C Daniel Gelatt Jr, and Mario P Vecchi. "Optimization by simulated annealing." In: *science* 220.4598 (1983), pp. 671–680.

[136] Knit It Now LLC. *KnitItNow Pattern Library*. [Online]. Available from: https://www.knititnow.com/knit/catalog.cfm. 2013.

[137] F. B. Kong, X. G. Ming, L. Wang, X. H. Wang, and P. P. Wang. "On Modular Products Development." In: *Concurrent Eng.* 17.4 (2009), pp. 291–300. ISSN: 1063-293X. DOI: 10.1177/1063293X09353974.

[138] Yuki Koyama, Issei Sato, and Masataka Goto. "Sequential Gallery for Interactive Visual Design Optimization." In: *ACM Trans. Graph.* 39.4 (2020). ISSN: 0730-0301. DOI: 10.1145/3386569.3392444. URL: https://doi.org/10.1145/3386569.3392444.

[139] Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. "Sequential Line Search for Efficient Visual Design Optimization by Crowds." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073598. URL: https://doi.org/10.1145/3072959.3073598.

[140] Yuki Koyama, Shinjiro Sueda, Emma Steinhardt, Takeo Igarashi, Ariel Shamir, and Wojciech Matusik. "AutoConnect: Computational Design of 3D-Printable Connectors." In: *ACM Trans. Graph.* 34.6 (2015). ISSN: 0730-0301. DOI: 10.1145/2816795.2818060. URL: https://doi.org/10.1145/2816795.2818060.

[141] Sivam Krish. "A practical generative design method." In: *Computer-Aided Design* 43.1 (2011), pp. 88–100. ISSN: 0010-4485. DOI: https://doi.org/10.1016/j.cad.2010.09.009. URL: https://www.sciencedirect.com/science/article/pii/S0010448510001764.

[142] Catherine J. Kudlick. "Disability History: Why We Need Another "Other"." In: *Am. Hist. Rev.* 108.3 (2003), pp. 763–793. ISSN: 0002-8762. DOI: 10.1086/ahr/108.3.763.

[143] Stacey Kuznetsov and Eric Paulos. "Rise of the Expert Amateur: DIY Projects, Communities, and Cultures." In: *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries*. NordiCHI '10. ACM, 2010, pp. 295–304. ISBN: 978-1-60558-934-3. DOI: 10.1145/1868914.1868950. URL: http://doi.acm.org/10.1145/1868914.1868950.

[144] Peter J. M. van Laarhoven and Emile H. L. Aarts. "Simulated annealing." In: *Simulated Annealing: Theory and Applications*. Dordrecht, The Netherlands: Springer, 1987, pp. 7–15. ISBN: 978-90-481-8438-5. DOI: 10.1007/978-94-015-7744-1_2.

[145] Richard E. Ladner. "Access and Empowerment: Commentary on "Computers and People with Disabilities"." In: *ACM Trans. Accessible Comput.* 1.2 (2008), pp. 1–5. ISSN: 1936-7228. DOI: 10.1145/1408760.1408765.

[146] Udaya Lakshmi, Megan Hofmann, Kelly Mack, Scott E. Hudson, Jennifer Mankoff, and Rosa I. Arriaga. "Medical Maker Response to COVID-19: Distributed Manufacturing Infrastructure for Stopgap Protective Equipment." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. URL: https://doi.org/10.1145/3411764.3445395.

[147] Manfred Lau, Akira Ohgawara, Jun Mitani, and Takeo Igarashi. "Converting 3D Furniture Models to Fabricatable Parts and Connectors." In: *ACM Trans. Graph.* 30.4 (2011). ISSN: 0730-0301. DOI: 10.1145/2010324.1964980. URL: https://doi.org/10.1145/2010324.1964980.

[148] Jonathan Leaf, Rundong Wu, Eston Schweickart, Doug L. James, and Steve Marschner. "Interactive Design of Periodic Yarn-level Cloth Patterns." In: *ACM Trans. Graph.* 37.6 (2018), 202:1–202:15. ISSN: 0730-0301. DOI: 10.1145/3272127.3275105. URL: http://doi.acm.org/10.1145/3272127.3275105.

[149] Joonhwan Lee, Jodi Forlizzi, and Scott E. Hudson. "Studying the Effectiveness of MOVE: A Contextually Optimized in-Vehicle Navigation System." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2005, 571–580. ISBN: 1581139985. URL: https://doi.org/10.1145/1054972.1055051.

[150] Joonhwan Lee, Jodi Forlizzi, Scott E. Hudson, and Soojin Jun. "Use of the Backseat Driving Technique in Evaluation of a Perceptually Optimized In-Car Navigation Display." In: *International Journal of Human–Computer Interaction* 31.2 (2015), pp. 128–138. DOI: 10.1080/10447318.2014.959104. eprint: https://doi.org/10.1080/10447318.2014.959104. URL: https://doi.org/10.1080/10447318.2014.959104.

[151] Dingzeyu Li, David I. W. Levin, Wojciech Matusik, and Changxi Zheng. "Acoustic Voxels: Computational Optimization of Modular Acoustic Filters." In: *ACM Trans. Graph.* 35.4 (2016). ISSN: 0730-0301. DOI: 10.1145/2897824.2925960. URL: https://doi.org/10.1145/2897824.2925960.

[152] Jiahao Li, Meilin Cui, Jeeeun Kim, and Xiang 'Anthony' Chen. "Romeo: A Design Tool for Embedding Transformable Parts in 3D Models to Robotically Augment Default Functionalities." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST '20. Association for Computing Machinery, 2020, 897–911. ISBN: 9781450375146. DOI: 10.1145/3379337.3415826. URL: https://doi.org/10.1145/3379337.3415826.

[153] Louis Li and Krzysztof Z. Gajos. "Adaptive Click-and-Cross: Adapting to Both Abilities and Task Improves Performance of Users with Impaired Dexterity." In: *Proceedings of the 19th International Conference on Intelligent User Interfaces*. IUI '14. Association for Computing Machinery, 2014, 299–304. ISBN: 9781450321846. DOI: 10.1145/2557500.2557511. URL: https://doi.org/10.1145/2557500.2557511.

[154] Jenny Lin, Vidya Narayanan, and James McCann. "Efficient Transfer Planning for Flat Knitting." In: *Proceedings of the 2Nd ACM Symposium on Computational Fabrication*. SCF '18. ACM, 2018, 1:1–1:7. ISBN: 978-1-4503-5854-5. DOI: 10.1145/3213512.3213515. URL: http://doi.acm.org/10.1145/3213512.3213515.

[155] Yvonna Lincoln and Egon Guba. "(1985)." In: *Naturalistic inquiry* (1992).

[156] Silvia Lindtner, Shaowen Bardzell, and Jeffrey Bardzell. "Reconstituting the Utopian Vision of Making: HCI After Technosolutionism." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. ACM, 2016, pp. 1390–1402. ISBN: 978-1-4503-3362-7. DOI: 10.1145/2858036.2858506. URL: http://doi.acm.org/10.1145/2858036.2858506.

[157] Silvia Lindtner, Garnet D. Hertz, and Paul Dourish. "Emerging Sites of HCI Innovation: Hackerspaces, Hardware Startups & Incubators." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '14. ACM, 2014, pp. 439–448. ISBN: 978-1-4503-2473-1. DOI: 10.1145/2556288.2557132. URL: http://doi.acm.org/10.1145/2556288.2557132.

[158] Helena R. Lourenço, Olivier C. Martin, and Thomas Stützle. "Iterated Local Search." In: *Handbook of Metaheuristics*. Boston, MA, USA: Springer, Boston, MA, 2003, pp. 320–353. ISBN: 978-1-4020-7263-5. DOI: 10.1007/0-306-48056-5_11.

[159] Lin Lu, Andrei Sharf, Haisen Zhao, Yuan Wei, Qingnan Fan, Xuelin Chen, Yann Savoye, Changhe Tu, Daniel Cohen-Or, and Baoquan Chen. "Build-to-Last: Strength to Weight 3D Printed Objects." In: *ACM Trans. Graph.* 33.4 (2014). ISSN: 0730-0301. DOI: 10.1145/2601097.2601168. URL: https://doi.org/10.1145/2601097.2601168.

[160] Yiyue Luo, Kui Wu, Tomás Palacios, and Wojciech Matusik. "KnitUI: Fabricating Interactive and Sensing Textiles with Machine Knitting." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445780. URL: https://doi.org/10.1145/3411764.3445780.

[161] Yiyue Luo, Kui Wu, Andrew Spielberg, Michael Foshey, Daniela Rus, Tomás Palacios, and Wojciech Matusik. "Digital Fabrication of Pneumatic Actuators with Integrated Sensing by Machine Knitting." In: *CHI Conference on Human Factors in Computing Systems*. CHI '22. New York, NY, USA: Association for Computing Machinery, 2022. ISBN: 9781450391573. DOI: 10.1145/3491102.3517577. URL: https://doi.org/10.1145/3491102.3517577.

[162] Allan MacLean, Kathleen Carter, Lennart Lövstrand, and Thomas Moran. "User-Tailorable Systems: Pressing the Issues with Buttons." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '90. Association for Computing Machinery, 1990, 175–182. ISBN: 0201509326. DOI: 10.1145/97243.97271. URL: https://doi.org/10.1145/97243.97271.

[163] Kelly Mack, Emma McDonnell, Venkatesh Potluri, Maggie Xu, Jailyn Zabala, Jeffrey Bigham, Jennifer Mankoff, and Cynthia Bennett. "Anticipate and Adjust: Cultivating Access in Human-Centered Methods." In: *CHI Conference on Human Factors in Computing Systems*. CHI '22. New York, NY, USA: Association for Computing Machinery, 2022. ISBN: 9781450391573. DOI: 10.1145/3491102.3501882. URL: https://doi.org/10.1145/3491102.3501882.

[164] Rachel P Maines. *Hedonizing technologies: paths to pleasure in hobbies and leisure*. JHU Press, 2009. ISBN: 9780801897948.

[165] Jennifer Mankoff, Megan Hofmann, Xiang 'Anthony' Chen, Scott E. Hudson, Amy Hurst, and Jeeeun Kim. "Consumer-Grade Fabrication and Its Potential to Revolutionize Accessibility." In: *Commun. ACM* 62.10 (2019), 64–75. ISSN: 0001-0782. DOI: 10.1145/3339824. URL: https://doi.org/10.1145/3339824.

[166] *Maple Global Optimization Toolbox Powered by Optimus - Optimization Software - Maplesoft*. [Online; accessed 1. Apr. 2021]. 2021. URL: https://www.maplesoft.com/products/toolboxes/globaloptimization.

[167] David R. Marshall and Deborah A. McGrew. "Creativity and Innovation in Health Care: Opening a Hospital Makerspace." In: *Nurse Leader* 15.1 (2017), pp. 56 –58. ISSN: 1541-4612. DOI: https://doi.org/10.1016/j.mnl.2016.10.002. URL: http://www.sciencedirect.com/science/article/pii/S1541461216302233.

[168] Justin Matejka, Michael Glueck, Erin Bradner, Ali Hashemi, Tovi Grossman, and George Fitzmaurice. "Dream Lens: Exploration and Visualization of Large-Scale Generative Design Datasets." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. New York, NY, USA: Association for Computing Machinery, 2018, 1–12. ISBN: 9781450356206. DOI: 10.1145/3173574.3173943. URL: https://doi.org/10.1145/3173574.3173943.

[169] James McCann. *The "Knitout" (.k) File Format*. 2017. URL: https://textiles-lab.github.io/knitout/knitout.html.

[170] James McCann, Lea Albaugh, Vidya Narayanan, April Grow, Wojciech Matusik, Jennifer Mankoff, and Jessica Hodgins. "A Compiler for 3D Machine Knitting." In: *ACM Trans. Graph.* 35.4 (2016), 49:1–49:11. ISSN: 0730-0301. DOI: 10.1145/2897824.2925940. URL: http://doi.acm.org/10.1145/2897824.2925940.

[171] Samantha McDonald, Niara Comrie, Erin Buehler, Nicholas Carter, Braxton Dubin, Karen Gordes, Sandy McCombe-Waller, and Amy Hurst. "Uncovering Challenges and Opportunities for 3D Printing Assistive Technology with Physical Therapists." In: *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '16. ACM, 2016, pp. 131–139. ISBN: 978-1-4503-4124-0. DOI: 10.1145/2982142.2982162. URL: http://doi.acm.org/10.1145/2982142.2982162.

[172] Samantha McDonald, Niara Comrie, Erin Buehler, Nicholas Carter, Braxton Dubin, Karen Gordes, Sandy McCombe-Waller, and Amy Hurst. "Uncovering Challenges and Opportunities for 3D Printing Assistive Technology with Physical Therapists." In: *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '16. ACM, 2016, pp. 131–139. ISBN: 978-1-4503-4124-0. DOI: 10.1145/2982142.2982162. URL: http://doi.acm.org/10.1145/2982142.2982162.

[173] I. Medicine, B.P.H.P.H. Practice, C.P.H.E.F.D.A.C. Process, and T. Wizemann. *Public Health Effectiveness of the FDA 510(k) Clearance Process: Balancing Patient Safety and Innovation: Workshop Report*. National Academies Press, 2010. ISBN: 9780309158497. URL: https://books.google.com/books?id=1CNhAgAAQBAJ.

[174] Vittorio Megaro, Bernhard Thomaszewski, Maurizio Nitti, Otmar Hilliges, Markus Gross, and Stelian Coros. "Interactive Design of 3D-Printable Robotic Creatures." In: *ACM Trans. Graph.* 34.6 (2015). ISSN: 0730-0301. DOI: 10.1145/2816795.2818137. URL: https://doi.org/10.1145/2816795.2818137.

[175] Janis Lena Meissner, John Vines, Janice McLaughlin, Thomas Nappey, Jekaterina Maksimova, and Peter Wright. "Do-It-Yourself Empowerment As Experienced by Novice Makers with Disabilities." In: *Proceedings of the 2017 Conference on Designing Interactive Systems*. DIS '17. ACM, 2017, pp. 1053–1065. ISBN: 978-1-4503-4922-2. DOI: 10.1145/3064663.3064674. URL: http://doi.acm.org/10.1145/3064663.3064674.

[176] Michael Meißner and Bernd Eberhardt. "The art of knitted fabrics, realistic & physically based modelling of knitted patterns." In: *Computer Graphics Forum*. Vol. 17. 3. Wiley Online Library. 1998, pp. 355–362.

[177] Michael Menninger. *Three Chips for Sister marsha Transcript*. 2014. URL: http://www.goodeatsfanpage.com/season3/cookie/cookietranscript.htm.

[178] Joshua A Miele, Steven Landau, and Deborah Gilden. "Talking TMAP: Automated generation of audio-tactile maps using Smith-Kettlewell's TMAP software." In: *British Journal of Visual Impairment* 24.2 (2006), pp. 93–100.

[179] Matthew B Miles, A Michael Huberman, Michael A Huberman, and Michael Huberman. *Qualitative data analysis: An expanded sourcebook*. sage, 1994.

[180] N. Mladenović and P. Hansen. "Variable neighborhood search." In: *Computers & Operations Research* 24.11 (1997), pp. 1097–1100. ISSN: 0305-0548. DOI: 10.1016/S0305-0548(97)00031-2.

[181] Yuki Mori and Takeo Igarashi. "Plushie: An Interactive Design System for Plush Toys." In: *ACM SIGGRAPH 2007 Papers*. SIGGRAPH '07. ACM, 2007. DOI: 10.1145/1275808.1276433. URL: http://doi.acm.org/10.1145/1275808.1276433.

[182] Martez E. Mott and Jacob O. Wobbrock. "Cluster Touch: Improving Touch Accuracy on Smartphones for People with Motor and Situational Impairments." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. Association for Computing Machinery, 2019, 1–14. ISBN: 9781450359702. DOI: 10.1145/3290605.3300257. URL: https://doi.org/10.1145/3290605.3300257.

[183] Clare Mountain-Manipon. *How to Grade Knitting Patterns Using a Spreadsheet | Sister Mountain*. 2018. URL: https://www.sistermountain.com/blog/grade-knitting-patterns-spreadsheet.

[184] S. Mukhopadhyay and R. Poojary. "A review on 3D printing: Advancement in healthcare technology." In: *2018 Advances in Science and Engineering Technology International Conferences (ASET)*. 2018, pp. 1–5. DOI: 10.1109/ICASET.2018.8376890.

[185] Vidya Narayanan, Lea Albaugh, Jessica Hodgins, Stelian Coros, and James McCann. "Automatic Machine Knitting of 3D Meshes." In: *ACM Trans. Graph.* 37.3 (2018), 35:1–35:15. ISSN: 0730-0301. DOI: 10.1145/3186265. URL: http://doi.acm.org/10.1145/3186265.

[186] Vidya Narayanan, Kui Wu, Cem Yuksel, and James McCann. "Visual Knit Programming." In: *ACM Trans. Graph.* 38.4 (2019). ISSN: 0730-0301.

[187] Thomas Klaus Nindel, Tomáš Iser, Tobias Rittig, Alexander Wilkie, and Jaroslav Křivánek. "A Gradient-Based Framework for 3D Print Appearance Optimization." In: *ACM Trans. Graph.* 40.4 (2021). ISSN: 0730-0301. DOI: 10.1145/3450626.3459844. URL: https://doi.org/10.1145/3450626.3459844.

[188] D. Norman. *The Design of Everyday Things: Revised and Expanded Edition*. Basic Books, 2013. ISBN: 9780465072996. URL: https://books.google.com/books?id=nVQPAAAAQBAJ.

[189] Aisling Ann O'Kane, Yi Han, and Rosa I. Arriaga. "Varied & Bespoke Caregiver Needs: Organizing and Communicating Diabetes Care for Children in the DIY Era." In: *Proceedings of the 10th EAI International Conference on Pervasive Computing Technologies for Healthcare*. PervasiveHealth '16. ICST, Brussels, Belgium: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2016, pp. 9–12. ISBN: 978-1-63190-051-8. URL: http://dl.acm.org/citation.cfm?id=3021319.3021321.

[190] Lora Oehlberg, Wesley Willett, and Wendy E. Mackay. "Patterns of Physical Design Remixing in Online Maker Communities." In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. Association for Computing Machinery, 2015, 639–648. ISBN: 9781450331456. DOI: 10.1145/2702123.2702175. URL: https://doi.org/10.1145/2702123.2702175.

[191] Johanna Okerlund and David Wilson. "DIY Assistive Technology for Others: Considering Social Impacts and Opportunities to Leverage HCI Techniques." In: *Proceedings of FabLearn 2019*. FL2019. Association for Computing Machinery, 2019, 152–155. ISBN: 9781450362443. DOI: 10.1145/3311890.3311914. URL: https://doi.org/10.1145/3311890.3311914.

[192] Dan R. Olsen. "Evaluating User Interface Systems Research." In: *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology*. UIST '07. New York, NY, USA: Association for Computing Machinery, 2007, 251–258. ISBN: 9781595936790. DOI: 10.1145/1294211.1294256. URL: https://doi.org/10.1145/1294211.1294256.

[193] OpenStreetMap Community Contributors. *OpenStreetMap*. [Online; accessed 25. Aug. 2021]. OpenStreetMap, 2021. URL: https://www.openstreetmap.org/.

[194] *Opt4J*. [Online; accessed 1. Apr. 2021]. 2021. URL: https://sdarg.github.io/opt4j.

[195] *Optimization Algorithm Toolkit (OAT)*. [Online; accessed 1. Apr. 2021]. 2021. URL: https://www.onworks.net/software/app-optimization-algorithm-toolkit-oat.

[196] *Optimization Toolbox*. [Online; accessed 1. Apr. 2021]. 2021. URL: https://www.mathworks.com/products/optimization.html.

[197] Jeremiah Parry-Hill, Patrick C. Shih, Jennifer Mankoff, and Daniel Ashbrook. "Understanding Volunteer AT Fabricators: Opportunities and Challenges in DIY-AT for Others in e-NABLE." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, Denver, CO, USA, May 06-11, 2017*. 2017, pp. 6184–6194. DOI: 10.1145/3025453.3026045. URL: http://doi.acm.org/10.1145/3025453.3026045.

[198] Jeremiah Parry-Hill, Patrick C Shih, Jennifer Mankoff, and Daniel Ashbrook. "Understanding Volunteer AT Fabricators: Opportunities and Challenges in DIY-AT for Others in e-NABLE." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM. 2017, pp. 6184–6194.

[199] Helen Petrie, Valerie Johnson, Thomas Strothotte, Andreas Raab, Rainer Michel, Lars Reichert, and Axel Schalt. "MoBIC: An Aid to Increase the Independent Mobility of Blind Travellers." In: *British Journal of Visual Impairment* 15.2 (1997), pp. 63–66. DOI: 10.1177/026461969701500205. eprint: https://doi.org/10.1177/026461969701500205. URL: https://doi.org/10.1177/026461969701500205.

[200] Betsy Phillips and Hongxin Zhao. "Predictors of Assistive Technology Abandonment." In: *Assistive Technol.* 5.1 (1993), pp. 36–45. ISSN: 1040-0435. DOI: 10.1080/10400435.1993.10132205.

[201] Michal Piovarči, David I. W. Levin, Jason Rebello, Desai Chen, Roman Ďurikovič, Hanspeter Pfister, Wojciech Matusik, and Piotr Didyk. "An Interaction-Aware, Perceptual Model for Non-Linear Elastic Objects." In: *ACM Trans. Graph.* 35.4 (2016). ISSN: 0730-0301. DOI: 10.1145/2897824.2925885. URL: https://doi.org/10.1145/2897824.2925885.

[202] Mariana Popescu, Matthias Rippmann, Tom Van Mele, and Philippe Block. "Automated generation of knit patterns for non-developable surfaces." In: *Humanizing Digital Reality*. Springer, 2018, pp. 271–284.

[203] Romain Prévost, Emily Whiting, Sylvain Lefebvre, and Olga Sorkine-Hornung. "Make It Stand: Balancing Shapes for 3D Fabrication." In: *ACM Trans. Graph.* 32.4 (2013). ISSN: 0730-0301. DOI: 10.1145/2461912.2461957. URL: https://doi.org/10.1145/2461912.2461957.

[204] Halley P. Profita, Abigale Stangl, Laura Matuszewska, Sigrunn Sky, and Shaun K. Kane. "Nothing to Hide: Aesthetic Customization of Hearing Aids and Cochlear Implants in an Online Community." In: *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '16. ACM, 2016, pp. 219–227. ISBN: 978-1-4503-4124-0. DOI: 10.1145/2982142.2982159. URL: http://doi.acm.org/10.1145/2982142.2982159.

[205] Halley P. Profita, Abigale Stangl, Laura Matuszewska, Sigrunn Sky, Raja Kushalnagar, and Shaun K. Kane. ""Wear It Loud": How and Why Hearing Aid and Cochlear Implant Users Customize Their Devices." In: *ACM Trans. Access. Comput.* 11.3 (2018), 13:1–13:32. ISSN: 1936-7228. DOI: 10.1145/3214382. URL: http://doi.acm.org/10.1145/3214382.

[206] *Prusa Protective Face Shield RC3 US version | NIH 3D Print Exchange*. 2020. URL: https://3dprint.nih.gov/discover/3DPX-013409.

[207] Christian Quintero. "A review: accessible technology through participatory design." In: *Disability and Rehabilitation: Assistive Technology* (2020), pp. 1–7. ISSN: 1748-3107. DOI: 10.1080/17483107.2020.1785564.

[208] *Ravelry: Home*. 2017. URL: https://www.ravelry.com/.

[209] Timothy V Roberts, Chris Hodge, Gerard Sutton, Michael Lawless, and contributors to the Vision Eye Institute IOL outcomes registry. "Comparison of Hill-radial basis function, Barrett Universal and current third generation formulas for the calculation of intraocular lens power during cataract surgery." In: *Clinical & Experimental Ophthalmology* 46.3 (2018), pp. 240–246. DOI: https://doi.org/10.1111/ceo.13034. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ceo.13034. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/ceo.13034.

[210] Daniela K. Rosner, Silvia Lindtner, Ingrid Erickson, Laura Forlano, Steven J. Jackson, and Beth Kolko. "Making Cultures: Building Things & Building Communities." In: *Proceedings of the Companion Publication of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*. CSCW Companion '14. ACM, 2014, pp. 113–116. ISBN: 978-1-4503-2541-7. DOI: 10.1145/2556420.2556852. URL: http://doi.acm.org/10.1145/2556420.2556852.

[211] Daniela Rosner and Jonathan Bean. "Learning from IKEA Hacking: I'm Not One to Decoupage a Tabletop and Call It a Day." In: CHI '09. Association for Computing Machinery, 2009, 419–422. ISBN: 9781605582467. DOI: 10.1145/1518701.1518768. URL: https://doi.org/10.1145/1518701.1518768.

[212] Thijs Jan Roumen, Willi Müller, and Patrick Baudisch. "Grafter: Remixing 3D-Printed Machines." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. ACM, 2018, 63:1–63:12. ISBN: 978-1-4503-5620-6. DOI: 10.1145/3173574.3173637. URL: http://doi.acm.org/10.1145/3173574.3173637.

[213] Thijs Roumen. "Portable Laser Cutting." In: *Adjunct Publication of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST '20 Adjunct. New York, NY, USA: Association for Computing Machinery, 2020, 157–161. ISBN: 9781450375153. DOI: 10.1145/3379350.3415802. URL: https://doi.org/10.1145/3379350.3415802.

[214] *SCF '21: Symposium on Computational Fabrication*. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450390903.

[215] Robert Savage. "Fusion 360 Introduction to Generative Design - Autodesk University - Medium." In: *Medium* (2020). ISSN: 4244-2055. URL: https://medium.com/autodesk-university/fusion-360-introduction-to-generative-design-eff424b42055.

[216] Douglas Schuler and Aki Namioka. *Participatory design: Principles and practices*. CRC Press, 1993. ISBN: 9780805809510.

[217] Adriana Schulz. "Computational Design for the Next Manufacturing Revolution." PhD thesis. 2018.

[218] Adriana Schulz, Ariel Shamir, David I. W. Levin, Pitchaya Sitthi-amorn, and Wojciech Matusik. "Design and Fabrication by Example." In: 33.4 (2014). ISSN: 0730-0301. DOI: 10.1145/2601097.2601127. URL: https://doi.org/10.1145/2601097.2601127.

[219] Adriana Schulz, Jie Xu, Bo Zhu, Changxi Zheng, Eitan Grinspun, and Wojciech Matusik. "Interactive Design Space Exploration and Optimization for CAD Models." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073688. URL: https://doi.org/10.1145/3072959.3073688.

[220] Lei Shi. "Talkabel: A Labeling Method for 3D Printed Models." In: *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility*. ASSETS '15. Association for Computing Machinery, 2015, 361–362. ISBN: 9781450334006. DOI: 10.1145/2700648.2811327. URL: https://doi.org/10.1145/2700648.2811327.

[221] Lei Shi, Idan Zelzer, Catherine Feng, and Shiri Azenkot. "Tickers and Talker: An Accessible Labeling Toolkit for 3D Printed Models." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. Association for Computing Machinery, 2016, 4896–4907. ISBN: 9781450333627. DOI: 10.1145/2858036.2858507. URL: https://doi.org/10.1145/2858036.2858507.

[222] Lei Shi, Yuhang Zhao, and Shiri Azenkot. "Designing Interactions for 3D Printed Models with Blind People." In: *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '17. ACM, 2017, pp. 200–209. ISBN: 978-1-4503-4926-0. DOI: 10.1145/3132525.3132549. URL: http://doi.acm.org/10.1145/3132525.3132549.

[223] Lei Shi, Yuhang Zhao, and Shiri Azenkot. "Markit and Talkit: A Low-Barrier Toolkit to Augment 3D Printed Models with Audio Annotations." In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. UIST '17. Québec City, QC, Canada: Association for Computing Machinery, 2017, 493–506. ISBN: 9781450349819. DOI: 10.1145/3126594.3126650. URL: https://doi.org/10.1145/3126594.3126650.

[224] Shima Seiki. *SDS-ONE Apex3*. [Online]. Available from: http://www.shimaseiki.com/product/design/sdsone_apex/flat/. 2011.

[225] Evan Shimizu, Matthew Fisher, Sylvain Paris, James McCann, and Kayvon Fatahalian. "Design Adjectives: A Framework for Interactive Model-Guided Exploration of Parameterized Design Spaces." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST '20. Association for Computing Machinery, 2020, 261–278. ISBN: 9781450375146. DOI: 10.1145/3379337.3415866. URL: https://doi.org/10.1145/3379337.3415866.

[226] Evan Shimizu, Matthew Fisher, Sylvain Paris, James McCann, and Kayvon Fatahalian. "Design Adjectives: A Framework for Interactive Model-Guided Exploration of Parameterized Design Spaces." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST '20. New York, NY, USA: Association for Computing Machinery, 2020, 261–278. ISBN: 9781450375146. DOI: 10.1145/3379337.3415866. URL: https://doi.org/10.1145/3379337.3415866.

[227] Kristen Shinohara, Cynthia L. Bennett, Wanda Pratt, and Jacob O. Wobbrock. "Tenets for Social Accessibility: Towards Humanizing Disabled People in Design." In: 11.1 (2018). ISSN: 1936-7228. DOI: 10.1145/3178855. URL: https://doi.org/10.1145/3178855.

[228] Alexa F. Siu, Son Kim, Joshua A. Miele, and Sean Follmer. "ShapeCAD: An Accessible 3D Modelling Workflow for the Blind and Visually-Impaired Via 2.5D Shape Displays." In: *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '19. Association for Computing Machinery, 2019, 342–354. ISBN: 9781450366762. DOI: 10.1145/3308561.3353782. URL: https://doi.org/10.1145/3308561.3353782.

[229] Alexa F. Siu, Joshua Miele, and Sean Follmer. "An Accessible CAD Workflow Using Programming of 3D Models and Preview Rendering in A 2.5D Shape Display." In: *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '18. Association for Computing Machinery, 2018, 343–345. ISBN: 9781450356503. DOI: 10.1145/3234695.3240996. URL: https://doi.org/10.1145/3234695.3240996.

[230] Mélina Skouras, Bernhard Thomaszewski, Peter Kaufmann, Akash Garg, Bernd Bickel, Eitan Grinspun, and Markus Gross. "Designing Inflatable Structures." In: *ACM Trans. Graph.* 33.4 (2014). ISSN: 0730-0301. DOI: 10.1145/2601097.2601166. URL: https://doi.org/10.1145/2601097.2601166.

[231] Karin Slegers, Kristel Kouwenberg, Tereza Loučova, and Ramon Daniels. "Makers in Healthcare: The Role of Occupational Therapists in the Design of DIY Assistive Technology." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20. Association for Computing Machinery, 2020, 1–11. ISBN: 9781450367080. DOI: 10.1145/3313831.3376685. URL: https://doi.org/10.1145/3313831.3376685.

[232] Douglas G Smith and Ernest M Burgess. "The use of CAD/CAM technology in prosthetics and orthotics-current clinical models and a view to the future." In: *Journal of rehabilitation research and development* 38.3 (2001), pp. 327–334.

[233] Noah Smith. *Coronavirus In Colorado: Rural Colorado Hospital Receives Much-Needed PPE*. en-US. Denver, Colorado, USA, 2020. URL: https://denver.cbslocal.com/2020/04/04/coronavirus-holyoke-colorado-ppe/.

[234] Soft Byte LTD. *DesignaKnit 8*. [Online]. Available from: https://softbyte.co.uk/designaknit.htm. 2012.

[235] Norbert Sparrow. *FDA Tackles Opportunities, Challenges of 3D-printed Medical Devices*. en. Newsletter. 2014. URL: https://www.plasticstoday.com/content/fda-tackles-opportunities-challenges-3d-printed-medical-devices/13081585320639.

[236] Lesley Stanfield and Melody Griffiths. *The Essential Stitch Collection*. English. The Reader's Digest Association, Inc, 2010. ISBN: 978-1-60652-043-7.

[237] Abigale Stangl, Ann Cunningham, Lou Ann Blake, and Tom Yeh. "Defining Problems of Practices to Advance Inclusive Tactile Media Consumption and Production." In: *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '19. Association for Computing Machinery, 2019, 329–341. ISBN: 9781450366762. DOI: 10.1145/3308561.3353778. URL: https://doi.org/10.1145/3308561.3353778.

[238] Abigale Stangl, Chia-Lo Hsu, and Tom Yeh. "Transcribing Across the Senses: Community Efforts to Create 3D Printable Accessible Tactile Pictures for Young Children with Visual Impairments." In: *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility*. ASSETS '15. Association for Computing Machinery, 2015, 127–137. ISBN: 9781450334006. DOI: 10.1145/2700648.2809854. URL: https://doi.org/10.1145/2700648.2809854.

[239] Ondrej Stava, Juraj Vanek, Bedrich Benes, Nathan Carr, and Radomír Měch. "Stress Relief: Improving Structural Strength of 3D Printable Objects." In: *ACM Trans. Graph.* 31.4 (2012). ISSN: 0730-0301. DOI: 10.1145/2185520.2185544. URL: https://doi.org/10.1145/2185520.2185544.

[240] Stoll. *M1Plus pattern software*. [Online]. Available from: http://www.stoll.com/stoll_software_solutions_en_4/pattern_software_m1plus/3_1. 2011.

[241] Cristiano Storni. "A Personal Perspective on Research Through Design." In: *interactions* 22.4 (2015), pp. 74–76. ISSN: 1072-5520. DOI: 10.1145/2786974. URL: http://doi.acm.org/10.1145/2786974.

[242] T. Stützle. "Local search algorithms for combinatorial problems - analysis, improvements, and new applications." In: *DISKI*. 1999.

[243] Mike Szczys and Elliot Williams. *Hackaday*. en. 2013. URL: https://hackaday.com/about/.

[244] Joshua G. Tanenbaum, Amanda M. Williams, Audrey Desjardins, and Karen Tanenbaum. "Democratizing Technology: Pleasure, Utility and Expressiveness in DIY and Maker Practice." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '13. ACM, 2013, pp. 2603–2612. ISBN: 978-1-4503-1899-0. DOI: 10.1145/2470654.2481360. URL: http://doi.acm.org/10.1145/2470654.2481360.

[245] Joshua G. Tanenbaum, Amanda M. Williams, Audrey Desjardins, and Karen Tanenbaum. "Democratizing Technology: Pleasure, Utility and Expressiveness in DIY and Maker Practice." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '13. ACM, 2013, pp. 2603–2612. ISBN: 978-1-4503-1899-0. DOI: 10.1145/2470654.2481360. URL: http://doi.acm.org/10.1145/2470654.2481360.

[246] Brandon T. Taylor, Anind K. Dey, Dan P. Siewiorek, and Asim Smailagic. "TactileMaps.Net: A Web Interface for Generating Customized 3D-Printable Tactile Maps." In: *Proceedings of the 17th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '15. New York, NY, USA: Association for Computing Machinery, 2015, 427–428. ISBN: 9781450334006. DOI: 10.1145/2700648.2811336. URL: https://doi.org/10.1145/2700648.2811336.

[247] Brandon Taylor, Anind Dey, Dan Siewiorek, and Asim Smailagic. "Customizable 3D Printed Tactile Maps as Interactive Overlays." In: *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '16. Association for Computing Machinery, 2016, 71–79. ISBN: 9781450341240. DOI: 10.1145/2982142.2982167. URL: https://doi.org/10.1145/2982142.2982167.

[248] Brandon Taylor, Anind Dey, Dan Siewiorek, and Asim Smailagic. "Customizable 3D Printed Tactile Maps as Interactive Overlays." In: *Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '16. New York, NY, USA: Association for Computing Machinery, 2016, 71–79. ISBN: 9781450341240. DOI: 10.1145/2982142.2982167. URL: https://doi.org/10.1145/2982142.2982167.

[249] *TensorFlow Model Optimization*. [Online; accessed 1. Apr. 2021]. 2020. URL: https://www.tensorflow.org/model_optimization.

[250] Austin L. Toombs, Shaowen Bardzell, and Jeffrey Bardzell. "The Proper Care and Feeding of Hackerspaces: Care Ethics and Cultures of Making." In: *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. ACM, 2015, pp. 629–638. ISBN: 978-1-4503-3145-6. DOI: 10.1145/2702123.2702522. URL: http://doi.acm.org/10.1145/2702123.2702522.

[251] *UIST '21: The 34th Annual ACM Symposium on User Interface Software and Technology*. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450386357.

[252] Nobuyuki Umetani, Takeo Igarashi, and Niloy J. Mitra. "Guided Exploration of Physically Valid Shapes for Furniture Design." In: *ACM Trans. Graph.* 31.4 (2012). ISSN: 0730-0301. DOI: 10.1145/2185520.2185582. URL: https://doi.org/10.1145/2185520.2185582.

[253] Nobuyuki Umetani, Yuki Koyama, Ryan Schmidt, and Takeo Igarashi. "Pteromys: Interactive Design and Optimization of Free-Formed Free-Flight Model Airplanes." In: *ACM Trans. Graph.* 33.4 (2014). ISSN: 0730-0301. DOI: 10.1145/2601097.2601129. URL: https://doi.org/10.1145/2601097.2601129.

[254] J. Vanek, J. A. G. Galicia, and B. Benes. "Clever Support: Efficient Support Structure Generation for Digital Fabrication." In: *Comput. Graphics Forum* 33.5 (2014), pp. 117–125. ISSN: 0167-7055. DOI: 10.1111/cgf.12437.

[255] C. Lee Ventola. "Medical Applications for 3D Printing: Current and Projected Uses." In: *Pharmacy and Therapeutics* 39.10 (2014), pp. 704–711. ISSN: 1052-1372. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4189697/.

[256] Christos Voudouris and Edward P. K. Tsang. "Guided Local Search." In: *Handbook of Metaheuristics*. Boston, MA, USA: Springer, Boston, MA, 2003, pp. 185–218. ISBN: 978-1-4020-7263-5. DOI: 10.1007/0-306-48056-5_7.

[257] Dhaval Vyas. "Altruism and Wellbeing as Care Work in a Craft-Based Maker Culture." In: *Proc. ACM Hum.-Comput. Interact.* 3.GROUP (2019). DOI: 10.1145/3361120. URL: https://doi.org/10.1145/3361120.

[258] Jakub Wabiński and Albina Mościcka. "Automatic (tactile) map generation—a systematic literature review." In: *ISPRS International Journal of Geo-Information* 8.7 (2019), p. 293.

[259] Zheshen Wang, Baoxin Li, Terri Hedgpeth, and Teresa Haven. "Instant Tactile-Audio Map: Enabling Access to Digital Maps for People with Visual Impairment." In: *Proceedings of the 11th International ACM SIGACCESS Conference on Computers and Accessibility*. Assets '09. New York, NY, USA: Association for Computing Machinery, 2009, 43–50. ISBN: 9781605585581. DOI: 10.1145/1639642.1639652. URL: https://doi-org.offcampus.lib.washington.edu/10.1145/1639642.1639652.

[260] Zheshen Wang, Nan Li, and Baoxin Li. "Fast and independent access to map directions for people who are blind." In: *Interacting with computers* 24.2 (2012), pp. 91–106.

[261] *Welcome to Nightscout*. en. 2020. URL: http://www.nightscout.info/.

[262] *What is HIPAA*. URL: http://www.dhcs.ca.gov/formsandpubs/laws/hipaa/Pages/1.00WhatisHIPAA.aspx.

[263] *Where work happens*. en-US. 2020. URL: https://slack.com/.

[264] Emily Whiting, Nada Ouf, Liane Makatura, Christos Mousas, Zhenyu Shu, and Ladislav Kavan. "Environment-Scale Fabrication: Replicating Outdoor Climbing Experiences." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. New York, NY, USA: Association for Computing Machinery, 2017, 1794–1804. ISBN: 9781450346559. DOI: 10.1145/3025453.3025465. URL: https://doi.org/10.1145/3025453.3025465.

[265] Emily Whiting, Hijung Shin, Robert Wang, John Ochsendorf, and Frédo Durand. "Structural Optimization of 3D Masonry Buildings." In: *ACM Trans. Graph.* 31.6 (2012). ISSN: 0730-0301. DOI: 10.1145/2366145.2366178. URL: https://doi.org/10.1145/2366145.2366178.

[266] Michele A. Williams, Amy Hurst, and Shaun K. Kane. ""Pray before You Step out": Describing Personal and Situational Blind Navigation Behaviors." In: *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '13. New York, NY, USA: Association for Computing Machinery, 2013. ISBN: 9781450324052. DOI: 10.1145/2513383.2513449. URL: https://doi-org.offcampus.lib.washington.edu/10.1145/2513383.2513449.

[267] Sue Williams, Gini Laurie, and Flora Seielstad. *disability history museum–The Toomeyville Jr. Gazette Appears*. 1958. URL: https://www.disabilitymuseum.org/dhm/lib/detail.html?id=1278.

[268] Jacob O. Wobbrock, Shaun K. Kane, Krzysztof Z. Gajos, Susumu Harada, and Jon Froehlich. "Ability-Based Design: Concept, Principles and Examples." In: *ACM Trans. Access. Comput.* 3.3 (2011). ISSN: 1936-7228. DOI: 10.1145/1952383.1952384. URL: https://doi.org/10.1145/1952383.1952384.

[269] Jacob O. Wobbrock, Shaun K. Kane, Krzysztof Z. Gajos, Susumu Harada, and Jon Froehlich. "Ability-Based Design: Concept, Principles and Examples." In: *ACM Trans. Access. Comput.* 3.3 (2011). ISSN: 1936-7228. DOI: 10.1145/1952383.1952384. URL: https://doi.org/10.1145/1952383.1952384.

[270] Ruth Grace Wong. "On the Cutting Edge of Technical Knit Manufacturing." In: *Medium* (Dec. 2021). URL: https://medium.com/supplyframe-hardware/on-the-cutting-edge-of-technical-knit-manufacturing-72a17e68c7fc.

[271] Kui Wu, Hannah Swan, and Cem Yuksel. "Knittable Stitch Meshes." In: *ACM Trans. Graph.* 38.1 (2019), 10:1–10:13. ISSN: 0730-0301. DOI: 10.1145/3292481. URL: http://doi.acm.org/10.1145/3292481.

[272] Nur Yildirim, James McCann, and John Zimmerman. "Digital Fabrication Tools at Work: Probing Professionals' Current Needs and Desired Futures." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2020, 1–13. ISBN: 9781450367080. URL: https://doi.org/10.1145/3313831.3376621.

[273] Anna Young and Jose Gomez-Marquez. *Maker Health*. 2008. URL: http://www.makerhealth.co/.

[274] Anna Young and Jose Gomez-Marquez. *MakerNurse*. 2008. URL: http://makernurse.com/.

[275] Mehmet Ersin Yumer, Siddhartha Chaudhuri, Jessica K. Hodgins, and Levent Burak Kara. "Semantic Shape Editing Using Deformation Handles." In: *ACM Trans. Graph.* 34.4 (2015). ISSN: 0730-0301. DOI: 10.1145/2766908. URL: https://doi.org/10.1145/2766908.

[276] Ran Zhang, Thomas Auzinger, Duygu Ceylan, Wilmot Li, and Bernd Bickel. "Functionality-Aware Retargeting of Mechanisms to 3D Shapes." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073710. URL: https://doi.org/10.1145/3072959.3073710.

[277] Xiaoyi Zhang, Tracy Tran, Yuqian Sun, Ian Culhane, Shobhit Jain, James Fogarty, and Jennifer Mankoff. "Interactiles: 3D Printed Tactile Interfaces to Enhance Mobile Touchscreen Accessibility." In: *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility*. ASSETS '18. Association for Computing Machinery, 2018, 131–142. ISBN: 9781450356503. DOI: 10.1145/3234695.3236349. URL: https://doi.org/10.1145/3234695.3236349.

[278] Bo Zhu, Mélina Skouras, Desai Chen, and Wojciech Matusik. "Two-Scale Topology Optimization with Microstructures." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3095815. URL: https://doi.org/10.1145/3072959.3095815.

[279] Lifeng Zhu, Weiwei Xu, John Snyder, Yang Liu, Guoping Wang, and Baining Guo. "Motion-Guided Mechanical Toy Modeling." In: *ACM Trans. Graph.* 31.6 (2012). ISSN: 0730-0301. DOI: 10.1145/2366145.2366146. URL: https://doi.org/10.1145/2366145.2366146.

[280] John Zimmerman and Jodi Forlizzi. "The Role of Design Artifacts in Design Theory Construction." In: *Artifact* 2.1 (2008), pp. 41–45. ISSN: 1749-3463. DOI: 10.1080/17493460802276893.

[281] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. "Research through Design as a Method for Interaction Design Research in HCI." In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '07. Association for Computing Machinery, 2007, 493–502. ISBN: 9781595935939. DOI: 10.1145/1240624.1240704. URL: https://doi.org/10.1145/1240624.1240704.

[282] John Zimmerman, Erik Stolterman, and Jodi Forlizzi. "An Analysis and Critique of Research Through Design: Towards a Formalization of a Research Approach." In: *Proceedings of the 8th ACM Conference on Designing Interactive Systems*. DIS '10. ACM, 2010, pp. 310–319. ISBN: 978-1-4503-0103-9. DOI: 10.1145/1858171.1858228. URL: http://doi.acm.org/10.1145/1858171.1858228.

[283] Jorge M Zuniga, Adam M Carson, Jean M Peck, Thomas Kalina, Rakesh M Srivastava, and Kirk Peck. "The development of a low-cost three-dimensional printed shoulder, arm, and hand prostheses for children." In: *Prosthetics and orthotics international* 41.2 (2017), pp. 205–209.

GENERATIVE DESIGN AND OPTIMIZATION LITERATURE: 2016-2021

[1] Muhammad Abdullah, Martin Taraz, Yannis Kommana, Shohei Katakura, Robert Kovacs, Jotaro Shigeyama, Thijs Roumen, and Patrick Baudisch. "FastForce: Real-Time Reinforcement of Laser-Cut Structures." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445466. URL: https://doi.org/10.1145/3411764.3445466.

[2] Rahul Arora, Alec Jacobson, Timothy R. Langlois, Yijiang Huang, Caitlin Mueller, Wojciech Matusik, Ariel Shamir, Karan Singh, and David I. W. Levin. "Volumetric Michell Trusses for Parametric Design & Fabrication." In: *Proceedings of the ACM Symposium on Computational Fabrication*. SCF '19. New York, NY, USA: Association for Computing Machinery, 2019. ISBN: 9781450367950. DOI: 10.1145/3328939.3328999. URL: https://doi.org/10.1145/3328939.3328999.

[3] Thomas Auzinger, Wolfgang Heidrich, and Bernd Bickel. "Computational Design of Nanostructural Color for Additive Manufacturing." In: *ACM Trans. Graph.* 37.4 (2018). ISSN: 0730-0301. DOI: 10.1145/3197517.3201376. URL: https://doi.org/10.1145/3197517.3201376.

[4] Vahid Babaei, Kiril Vidimče, Michael Foshey, Alexandre Kaspar, Piotr Didyk, and Wojciech Matusik. "Color Contoning for 3D Printing." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073605. URL: https://doi.org/10.1145/3072959.3073605.

[5] Moritz Bächer, Benjamin Hepp, Fabrizio Pece, Paul G. Kry, Bernd Bickel, Bernhard Thomaszewski, and Otmar Hilliges. "DefSense: Computational Design of Customized Deformable Input Devices." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. New York, NY, USA: Association for Computing Machinery, 2016, 3806–3816. ISBN: 9781450333627. DOI: 10.1145/2858036.2858354. URL: https://doi.org/10.1145/2858036.2858354.

[6] Shajay Bhooshan, Tom Van Mele, and Philippe Block. "Morph & Slerp: Shape Description for 3D Printing of Concrete." In: *Symposium on Computational Fabrication*. SCF '20. New York, NY, USA: Association for Computing Machinery, 2020. ISBN: 9781450381703. DOI: 10.1145/3424630.3425413. URL: https://doi.org/10.1145/3424630.3425413.

[7] Ruei-Che Chang, Chih-An Tsao, Fang-Ying Liao, Seraphina Yong, Tom Yeh, and Bing-Yu Chen. "Daedalus in the Dark: Designing for Non-Visual Accessible Construction of Laser-Cut Architecture." In: *The 34th Annual ACM Symposium on User Interface Software and Technology*. UIST '21. New York, NY, USA: Association for Computing Machinery, 2021, 344–358. ISBN: 9781450386357. DOI: 10.1145/3472749.3474754. URL: https://doi.org/10.1145/3472749.3474754.

[8] Desai Chen, David I. W. Levin, Wojciech Matusik, and Danny M. Kaufman. "Dynamics-Aware Numerical Coarsening for Fabrication Design." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073669. URL: https://doi.org/10.1145/3072959.3073669.

[9] Xiang 'Anthony' Chen, Stelian Coros, and Scott E. Hudson. "Medley: A Library of Embeddables to Explore Rich Material Properties for 3D Printed Objects." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. New York, NY, USA: Association for Computing Machinery, 2018, 1–12. ISBN: 9781450356206. DOI: 10.1145/3173574.3173736. URL: https://doi.org/10.1145/3173574.3173736.

[10] Xiang 'Anthony' Chen, Stelian Coros, Jennifer Mankoff, and Scott E. Hudson. "Encore: 3D Printed Augmentation of Everyday Objects with Printed-Over, Affixed and Interlocked Attachments." In: *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. UIST '15. New York, NY, USA: Association for Computing Machinery, 2015, 73–82. ISBN: 9781450337793. DOI: 10.1145/2807442.2807498. URL: https://doi.org/10.1145/2807442.2807498.

[11] Xiang "Anthony" Chen, Ye Tao, Guanyun Wang, Runchang Kang, Tovi Grossman, Stelian Coros, and Scott E. Hudson. "Forte: User-Driven Generative Design." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Association for Computing Machinery, 2018. ISBN: 9781450356206. DOI: 10.1145/3173574.3174070. URL: https://doi.org/10.1145/3173574.3174070.

[12] Subramanian Chidambaram, Yunbo Zhang, Venkatraghavan Sundararajan, Niklas Elmqvist, and Karthik Ramani. "Shape Structuralizer: Design, Fabrication, and User-Driven Iterative Refinement of 3D Mesh Models." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. New York, NY, USA: Association for Computing Machinery, 2019, 1–12. ISBN: 9781450359702. DOI: 10.1145/3290605.3300893. URL: https://doi.org/10.1145/3290605.3300893.

[13] Sebastian Cucerca, Piotr Didyk, Hans-Peter Seidel, and Vahid Babaei. "Computational Image Marking on Metals via Laser Induced Heating." In: *ACM Trans. Graph.* 39.4 (2020). ISSN: 0730-0301. DOI: 10.1145/3386569.3392423. URL: https://doi.org/10.1145/3386569.3392423.

[14] Niraj Ramesh Dayama, Kashyap Todi, Taru Saarelainen, and Antti Oulasvirta. "GRIDS: Interactive Layout Design with Integer Programming." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2020, 1–13. ISBN: 9781450367080. URL: https://doi.org/10.1145/3313831.3376553.

[15] Ruta Desai, James McCann, and Stelian Coros. "Assembly-Aware Design of Printable Electromechanical Devices." In: *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST '18. New York, NY, USA: Association for Computing Machinery, 2018, 457–472. ISBN: 9781450359481. DOI: 10.1145/3242587.3242655. URL: https://doi.org/10.1145/3242587.3242655.

[16] Nikan Doosti, Julian Panetta, and Vahid Babaei. "Topology Optimization via Frequency Tuning of Neural Design Representations." In: *Symposium on Computational Fabrication*. SCF '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450390903. DOI: 10.1145/3485114.3485124. URL: https://doi.org/10.1145/3485114.3485124.

[17] Tao Du, Adriana Schulz, Bo Zhu, Bernd Bickel, and Wojciech Matusik. "Computational Multicopter Design." In: *ACM Trans. Graph.* 35.6 (2016). ISSN: 0730-0301. DOI: 10.1145/2980179.2982427. URL: https://doi.org/10.1145/2980179.2982427.

[18] Peitong Duan, Casimir Wierzynski, and Lama Nachman. "Optimizing User Interface Layouts via Gradient Descent." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2020, 1–12. ISBN: 9781450367080. URL: https://doi.org/10.1145/3313831.3376589.

[19] John J. Dudley, Jason T. Jacques, and Per Ola Kristensson. "Crowdsourcing Interface Feature Design with Bayesian Optimization." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. New York, NY, USA: Association for Computing Machinery, 2019, 1–12. ISBN: 9781450359702. DOI: 10.1145/3290605.3300482. URL: https://doi.org/10.1145/3290605.3300482.

[20] João Marcelo Evangelista Belo, Anna Maria Feit, Tiare Feuchtner, and Kaj Grønbæk. "XRgonomics: Facilitating the Creation of Ergonomic 3D Interfaces." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445349. URL: https://doi.org/10.1145/3411764.3445349.

[21] Christoph Gebhardt and Otmar Hilliges. "Optimization-Based User Support for Cinematographic Quadrotor Camera Target Framing." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445568. URL: https://doi.org/10.1145/3411764.3445568.

[22] Jun Gong, Olivia Seow, Cedric Honnet, Jack Forman, and Stefanie Mueller. "MetaSense: Integrating Sensing Capabilities into Mechanical Metamaterial." In: UIST '21. New York, NY, USA: Association for Computing Machinery, 2021, 1063–1073. ISBN: 9781450386357. DOI: 10.1145/3472749.3474806. URL: https://doi.org/10.1145/3472749.3474806.

[23] Jianzhe Gu, Vidya Narayanan, Guanyun Wang, Danli Luo, Harshika Jain, Kexin Lu, Fang Qin, Sijia Wang, James McCann, and Lining Yao. "Inverse Design Tool for Asymmetrical Self-Rising Surfaces with Color Texture." In: *Symposium on Computational Fabrication*. SCF '20. New York, NY, USA: Association for Computing Machinery, 2020. ISBN: 9781450381703. DOI: 10.1145/3424630.3425420. URL: https://doi.org/10.1145/3424630.3425420.

[24] Yue Hao, Yun-hyeong Kim, and Jyh-Ming Lien. "Synthesis of Fast and Collision-Free Folding of Polyhedral Nets." In: *Proceedings of the 2nd ACM Symposium on Computational Fabrication*. SCF '18. New York, NY, USA: Association for Computing Machinery, 2018. ISBN: 9781450358545. DOI: 10.1145/3213512.3213517. URL: https://doi.org/10.1145/3213512.3213517.

[25] Yue Hao and Jyh-Ming Lien. "Computational Laser Forming Origami of Convex Surfaces." In: *Proceedings of the ACM Symposium on Computational Fabrication*. SCF '19. New York, NY, USA: Association for Computing Machinery, 2019. ISBN: 9781450367950. DOI: 10.1145/3328939.3329006. URL: https://doi.org/10.1145/3328939.3329006.

[26] Jean Hergel, Kevin Hinz, Sylvain Lefebvre, and Bernhard Thomaszewski. "Extrusion-Based Ceramics Printing with Strictly-Continuous Deposition." In: *ACM Trans. Graph.* 38.6 (2019). ISSN: 0730-0301. DOI: 10.1145/3355089.3356509. URL: https://doi.org/10.1145/3355089.3356509.

[27] Alexandra Ion, David Lindlbauer, Philipp Herholz, Marc Alexa, and Patrick Baudisch. "Understanding Metamaterial Mechanisms." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. New York, NY, USA: Association for Computing Machinery, 2019, 1–14. ISBN: 9781450359702. DOI: 10.1145/3290605.3300877. URL: https://doi.org/10.1145/3290605.3300877.

[28] Yue Jiang, Wolfgang Stuerzlinger, Matthias Zwicker, and Christof Lutteroth. "ORCSolver: An Efficient Solver for Adaptive GUI Layout with OR-Constraints." In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2020, 1–14. ISBN: 9781450367080. URL: https://doi.org/10.1145/3313831.3376610.

[29] Yuhua Jin, Isabel Qamar, Michael Wessely, and Stefanie Mueller. "Photo- Chromeleon: Re- Programmable Multi-Color Textures Using Photochromic Dyes." In: *ACM SIGGRAPH 2020 Emerging Technologies*. SIGGRAPH '20. New York, NY, USA: Association for Computing Machinery, 2020. ISBN: 9781450379670. DOI: 10.1145/3388534.3407296. URL: https://doi.org/10.1145/3388534.3407296.

[30] Florian Kadner, Yannik Keller, and Constantin Rothkopf. "AdaptiFont: Increasing Individuals' Reading Speed with a Generative Font Model and Bayesian Optimization." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445140. URL: https://doi.org/10.1145/3411764.3445140.

[31] Martin Kilian, Aron Monszpart, and Niloy J. Mitra. "String Actuated Curved Folded Surfaces." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3015460. URL: https://doi.org/10.1145/3072959.3015460.

[32] Mina Konaković, Keenan Crane, Bailin Deng, Sofien Bouaziz, Daniel Piker, and Mark Pauly. "Beyond Developable: Computational Design and Fabrication with Auxetic Materials." In: *ACM Trans. Graph.* 35.4 (2016). ISSN: 0730-0301. DOI: 10.1145/2897824.2925944. URL: https://doi.org/10.1145/2897824.2925944.

[33] Yuki Koyama and Masataka Goto. "OptiMo: Optimization-Guided Motion Editing for Keyframe Character Animation." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. New York, NY, USA: Association for Computing Machinery, 2018, 1–12. ISBN: 9781450356206. DOI: 10.1145/3173574.3173735. URL: https://doi.org/10.1145/3173574.3173735.

[34] Yuki Koyama, Issei Sato, Daisuke Sakamoto, and Takeo Igarashi. "Sequential Line Search for Efficient Visual Design Optimization by Crowds." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073598. URL: https://doi.org/10.1145/3072959.3073598.

[35] Thomas Langerak, Juan José Zárate, Velko Vechev, David Lindlbauer, Daniele Panozzo, and Otmar Hilliges. "Optimal Control for Electromagnetic Haptic Guidance Systems." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. New York, NY, USA: Association for Computing Machinery, 2020, 951–965. ISBN: 9781450375146. URL: https://doi.org/10.1145/3379337.3415593.

[36] DoYoung Lee, Jiwan Kim, and Ian Oakley. "FingerText: Exploring and Optimizing Performance for Wearable, Mobile and One-Handed Typing." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445106. URL: https://doi.org/10.1145/3411764.3445106.

[37] Danny Leen, Tom Veuskens, Kris Luyten, and Raf Ramakers. "JigFab: Computational Fabrication of Constraints to Facilitate Woodworking with Power Tools." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. New York, NY, USA: Association for Computing Machinery, 2019, 1–12. ISBN: 9781450359702. DOI: 10.1145/3290605.3300386. URL: https://doi.org/10.1145/3290605.3300386.

[38] Dingzeyu Li, David I. W. Levin, Wojciech Matusik, and Changxi Zheng. "Acoustic Voxels: Computational Optimization of Modular Acoustic Filters." In: *ACM Trans. Graph.* 35.4 (2016). ISSN: 0730-0301. DOI: 10.1145/2897824.2925960. URL: https://doi.org/10.1145/2897824.2925960.

[39] David Lindlbauer, Anna Maria Feit, and Otmar Hilliges. "Context-Aware Online Adaptation of Mixed Reality Interfaces." In: *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. UIST '19. New York, NY, USA: Association for Computing Machinery, 2019, 147–160. ISBN: 9781450368162. DOI: 10.1145/3332165.3347945. URL: https://doi.org/10.1145/3332165.3347945.

[40] J. Derek Lomas, Jodi Forlizzi, Nikhil Poonwala, Nirmal Patel, Sharan Shodhan, Kishan Patel, Ken Koedinger, and Emma Brunskill. "Interface Design Optimization as a Multi-Armed Bandit Problem." In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. New York, NY, USA: Association for Computing Machinery, 2016, 4142–4153. ISBN: 9781450333627. DOI: 10.1145/2858036.2858425. URL: https://doi.org/10.1145/2858036.2858425.

[41] Granit Luzhnica and Eduardo Veas. "Optimising Encoding for Vibrotactile Skin Reading." In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. New York, NY, USA: Association for Computing Machinery, 2019, 1–14. ISBN: 9781450359702. DOI: 10.1145/3290605.3300465. URL: https://doi.org/10.1145/3290605.3300465.

[42] Jonàs Martínez, Haichuan Song, Jérémie Dumas, and Sylvain Lefebvre. "Orthotropic <i>k</i>-Nearest Foams for Additive Manufacturing." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073638. URL: https://doi.org/10.1145/3072959.3073638.

[43] Jess McIntosh, Hubert Dariusz Zajac, Andreea Nicoleta Stefan, Joanna Bergström, and Kasper Hornbæk. "Iteratively Adapting Avatars Using Task-Integrated Optimisation." In: *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. New York, NY, USA: Association for Computing Machinery, 2020, 709–721. ISBN: 9781450375146. URL: https://doi.org/10.1145/3379337.3415832.

[44] Vittorio Megaro, Jonas Zehnder, Moritz Bächer, Stelian Coros, Markus Gross, and Bernhard Thomaszewski. "A Computational Design Tool for Compliant Mechanisms." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073636. URL: https://doi.org/10.1145/3072959.3073636.

[45] Eder Miguel, Mathias Lepoutre, and Bernd Bickel. "Computational Design of Stable Planar-Rod Structures." In: *ACM Trans. Graph.* 35.4 (2016). ISSN: 0730-0301. DOI: 10.1145/2897824.2925978. URL: https://doi.org/10.1145/2897824.2925978.

[46] Roberto A. Montano Murillo, Sriram Subramanian, and Diego Martinez Plasencia. "Erg-O: Ergonomic Optimization of Immersive Virtual Environments." In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. UIST '17. New York, NY, USA: Association for Computing Machinery, 2017, 759–771. ISBN: 9781450349819. DOI: 10.1145/3126594.3126605. URL: https://doi.org/10.1145/3126594.3126605.

[47] Rafael Morales, Asier Marzo, Sriram Subramanian, and Diego Martínez. "LeviProps: Animating Levitated Optimized Fabric Structures Using Holographic Acoustic Tweezers." In: *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. UIST '19. New York, NY, USA: Association for Computing Machinery, 2019, 651–661. ISBN: 9781450368162. DOI: 10.1145/3332165.3347882. URL: https://doi.org/10.1145/3332165.3347882.

[48] Przemyslaw Musialski, Christian Hafner, Florian Rist, Michael Birsak, Michael Wimmer, and Leif Kobbelt. "Non-Linear Shape Optimization Using Local Subspace Projections." In: *ACM Trans. Graph.* 35.4 (2016). ISSN: 0730-0301. DOI: 10.1145/2897824.2925886. URL: https://doi.org/10.1145/2897824.2925886.

[49] Thomas Klaus Nindel, Tomáš Iser, Tobias Rittig, Alexander Wilkie, and Jaroslav Křivánek. "A Gradient-Based Framework for 3D Print Appearance Optimization." In: *ACM Trans. Graph.* 40.4 (2021). ISSN: 0730-0301. DOI: 10.1145/3450626.3459844. URL: https://doi.org/10.1145/3450626.3459844.

[50] Masa Ogata and Yuki Koyama. "A Computational Approach to Magnetic Force Feedback Design." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445631. URL: https://doi.org/10.1145/3411764.3445631.

[51] Julian Panetta, Florin Isvoranu, Tian Chen, Emmanuel Siéfert, Benoît Roman, and Mark Pauly. "Computational Inverse Design of Surface-Based Inflatables." In: *ACM Trans. Graph.* 40.4 (2021). ISSN: 0730-0301. DOI: 10.1145/3450626.3459789. URL: https://doi.org/10.1145/3450626.3459789.

[52] Seonwook Park et al. "AdaM: Adapting Multi-User Interfaces for Collaborative Environments in Real-Time." In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. New York, NY, USA: Association for Computing Machinery, 2018, 1–14. ISBN: 9781450356206. DOI: 10.1145/3173574.3173758. URL: https://doi.org/10.1145/3173574.3173758.

[53] Davide Pellis, Martin Kilian, Helmut Pottmann, and Mark Pauly. "Computational Design of Weingarten Surfaces." In: *ACM Trans. Graph.* 40.4 (2021). ISSN: 0730-0301. DOI: 10.1145/3450626.3459939. URL: https://doi.org/10.1145/3450626.3459939.

[54] Fang Qin, Huai-Yu Cheng, Rachel Sneeringer, Maria Vlachostergiou, Sampada Acharya, Haolin Liu, Carmel Majidi, Mohammad Islam, and Lining Yao. "ExoForm: Shape Memory and Self-Fusing Semi-Rigid Wearables." In: *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI EA '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380959. DOI: 10.1145/3411763.3451818. URL: https://doi.org/10.1145/3411763.3451818.

[55] Yingying Ren, Julian Panetta, Tian Chen, Florin Isvoranu, Samuel Poincloux, Christopher Brandt, Alison Martin, and Mark Pauly. "3D Weaving with Curved Ribbons." In: *ACM Trans. Graph.* 40.4 (2021). ISSN: 0730-0301. DOI: 10.1145/3450626.3459788. URL: https://doi.org/10.1145/3450626.3459788.

[56] Thijs Roumen et al. "Assembler3: 3D Reconstruction of Laser-Cut Models." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445453. URL: https://doi.org/10.1145/3411764.3445453.

[57] Ticha Sethapakdi, Daniel Anderson, Adrian Reginald Chua Sy, and Stefanie Mueller. "Fabricaide: Fabrication-Aware Design for 2D Cutting Machines." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445345. URL: https://doi.org/10.1145/3411764.3445345.

[58] Anna Shtengel, Roi Poranne, Olga Sorkine-Hornung, Shahar Z. Kovalsky, and Yaron Lipman. "Geometric Optimization via Composite Majorization." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073618. URL: https://doi.org/10.1145/3072959.3073618.

[59] Madlaina Signer, Alexandra Ion, and Olga Sorkine-Hornung. "Developable Metamaterials: Mass-Fabricable Metamaterials by Laser-Cutting Elastic Structures." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445666. URL: https://doi.org/10.1145/3411764.3445666.

[60] Haichuan Song, Jonàs Martínez, Pierre Bedell, Noémie Vennin, and Sylvain Lefebvre. "Colored Fused Filament Fabrication." In: *ACM Trans. Graph.* 38.5 (2019). ISSN: 0730-0301. DOI: 10.1145/3183793. URL: https://doi.org/10.1145/3183793.

[61] Denis Sumin, Tobias Rittig, Vahid Babaei, Thomas Nindel, Alexander Wilkie, Piotr Didyk, Bernd Bickel, Jaroslav Křivánek, Karol Myszkowski, and Tim Weyrich. "Geometry-Aware Scattering Compensation for 3D Printing." In: *ACM Trans. Graph.* 38.4 (2019). ISSN: 0730-0301. DOI: 10.1145/3306346.3322992. URL: https://doi.org/10.1145/3306346.3322992.

[62] Davi Colli Tozoni, Jérémie Dumas, Zhongshi Jiang, Julian Panetta, Daniele Panozzo, and Denis Zorin. "A Low-Parametric Rhombic Microstructure Family for Irregular Lattices." In: *ACM Trans. Graph.* 39.4 (2020). ISSN: 0730-0301. DOI: 10.1145/3386569.3392451. URL: https://doi.org/10.1145/3386569.3392451.

[63] Davi Colli Tozoni, Yunfan Zhou, and Denis Zorin. "Optimizing Contact-Based Assemblies." In: *ACM Trans. Graph.* 40.6 (2021). ISSN: 0730-0301. DOI: 10.1145/3478513.3480552. URL: https://doi.org/10.1145/3478513.3480552.

[64] Tom Valkeneers, Danny Leen, Daniel Ashbrook, and Raf Ramakers. "StackMold: Rapid Prototyping of Functional Multi-Material Objects with Selective Levels of Surface Details." In: *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. UIST '19. New York, NY, USA: Association for Computing Machinery, 2019, 687–699. ISBN: 9781450368162. DOI: 10.1145/3332165.3347915. URL: https://doi.org/10.1145/3332165.3347915.

[65] Michael Wessely, Yuhua Jin, Cattalyya Nuengsigkapian, Aleksei Kashapov, Isabel P. S. Qamar, Dzmitry Tsetserukou, and Stefanie Mueller. "ChromoUpdate: Fast Design Iteration of Photochromic Color Textures Using Grayscale Previews and Local Color Updates." In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445391. URL: https://doi.org/10.1145/3411764.3445391.

[66] Emily Whiting, Nada Ouf, Liane Makatura, Christos Mousas, Zhenyu Shu, and Ladislav Kavan. "Environment-Scale Fabrication: Replicating Outdoor Climbing Experiences." In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. New York, NY, USA: Association for Computing Machinery, 2017, 1794–1804. ISBN: 9781450346559. DOI: 10.1145/3025453.3025465. URL: https://doi.org/10.1145/3025453.3025465.

[67] Katja Wolff and Olga Sorkine-Hornung. "Wallpaper Pattern Alignment along Garment Seams." In: *ACM Trans. Graph.* 38.4 (2019). ISSN: 0730-0301. DOI: 10.1145/3306346.3322991. URL: https://doi.org/10.1145/3306346.3322991.

[68] Jiaxian Yao, Danny M. Kaufman, Yotam Gingold, and Maneesh Agrawala. "Interactive Design and Stability Analysis of Decorative Joinery for Furniture." In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3054740. URL: https://doi.org/10.1145/3072959.3054740.

[69] Christopher Yu, Keenan Crane, and Stelian Coros. “Computational Design of Telescoping Structures.” In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073673. URL: https://doi.org/10.1145/3072959.3073673.

[70] Ya-Ting Yue, Xiaolong Zhang, Yongliang Yang, Gang Ren, Yi-King Choi, and Wenping Wang. “WireDraw: 3D Wire Sculpturing Guided with Mixed Reality.” In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI ’17. New York, NY, USA: Association for Computing Machinery, 2017, 3693–3704. ISBN: 9781450346559. DOI: 10.1145/3025453.3025792. URL: https://doi.org/10.1145/3025453.3025792.

[71] Ran Zhang, Thomas Auzinger, Duygu Ceylan, Wilmot Li, and Bernd Bickel. “Functionality-Aware Retargeting of Mechanisms to 3D Shapes.” In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073710. URL: https://doi.org/10.1145/3072959.3073710.

[72] Ran Zhang, Thomas Auzinger, Duygu Ceylan, Wilmot Li, and Bernd Bickel. “Functionality-Aware Retargeting of Mechanisms to 3D Shapes.” In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3073710. URL: https://doi.org/10.1145/3072959.3073710.

[73] Xiaoting Zhang, Guoxin Fang, Chengkai Dai, Jouke Verlinden, Jun Wu, Emily Whiting, and Charlie C.L. Wang. “Thermal-Comfort Design of Personalized Casts.” In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. UIST ’17. New York, NY, USA: Association for Computing Machinery, 2017, 243–254. ISBN: 9781450349819. DOI: 10.1145/3126594.3126600. URL: https://doi.org/10.1145/3126594.3126600.

[74] Xiaoting Zhang, Guoxin Fang, Melina Skouras, Gwenda Gieseler, Charlie C. L. Wang, and Emily Whiting. “Computational Design of Fabric Formwork.” In: *ACM Trans. Graph.* 38.4 (2019). ISSN: 0730-0301. DOI: 10.1145/3306346.3322988. URL: https://doi.org/10.1145/3306346.3322988.

[75] Yongqi Zhang, Biao Xie, Haikun Huang, Elisa Ogawa, Tongjian You, and Lap-Fai Yu. “Pose-Guided Level Design.” In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI ’19. New York, NY, USA: Association for Computing Machinery, 2019, 1–12. ISBN: 9781450359702. DOI: 10.1145/3290605.3300784. URL: https://doi.org/10.1145/3290605.3300784.

[76] Bo Zhu, Mélina Skouras, Desai Chen, and Wojciech Matusik. “Two-Scale Topology Optimization with Microstructures.” In: *ACM Trans. Graph.* 36.4 (2017). ISSN: 0730-0301. DOI: 10.1145/3072959.3095815. URL: https://doi.org/10.1145/3072959.3095815.

[77] Amit Zoran and Dror Cohen. “Digital Konditorei: Programmable Taste Structures Using a Modular Mold.” In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI ’18. New York, NY, USA: Association for Computing Machinery, 2018, 1–9. ISBN: 9781450356206. DOI: 10.1145/3173574.3173974. URL: https://doi.org/10.1145/3173574.3173974.

COLOPHON

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". `classicthesis` is available for both LaTeX and LyX:

https://bitbucket.org/amiede/classicthesis/

Happy users of `classicthesis` usually send a real postcard to the author, a collection of postcards received so far is featured here:

http://postcards.miede.de/

Final Version as of July 25, 2022 (`classicthesis` version 4.2).