Sensory Encounters in the Age of Computation

by

Crystal Lee

A.B. with high honors (2015), A.M. (2016)
Stanford University

Submitted to the Program in Science, Technology, and Society in partial fulfillment of the requirements for the degree of Doctor of Philosophy in History, Anthropology, and Science, Technology, and Society at the Massachusetts Institute of Technology

September 2022

Author ……………………………………………………………………………………………………

History, Anthropology, and Science, Technology, and Society
July 28, 2022

Certified by……………………………………………………………………………………………

Graham M. Jones
Professor of Anthropology
Thesis Supervisor

Certified by……………………………………………………………………………………………

Arvind Satyanarayan
Associate Professor (without tenure) of Computer Science
Thesis Committee Member

Certified by……………………………………………………………………………………………………

William Deringer
Associate Professor of Science, Technology, and Society
Thesis Committee Member

Certified by……………………………………………………………………………………………………

Eden Medina
Associate Professor of Science, Technology, & Society (STS)
Thesis Committee Member

Accepted by……………………………………………………………………………………………………

Eden Medina
Associate Professor of Science, Technology, & Society (STS)
Director of Graduate Studies, History, Anthropology, and STS

Accepted by……………………………………………………………………………………………………

Deborah Fitzgerald
Leverett Howell and William King Cutten Professor of the History of Technology
Department Head, Program in Science, Technology, & Society

Sensory Encounters in the Age of Computation

Histories of Data Visualization and Human-Computer Interaction

by

Crystal Lee

Submitted to the Program in Science, Technology, and Society on July 5, 2022 in partial fulfillment of the requirements for the degree of Doctor of Philosophy in History, Anthropology, and Science, Technology, and Society at the Massachusetts Institute of Technology

Abstract

This dissertation investigates the process of curating, cleaning, visualizing, circulating, and manipulating data to understand the persuasive force of visual information in multimodal media. From the history of haptic interfaces to the data practices of social media communities across the US and China, this thesis uses historical and ethnographic methods to understand how users of quantitative information encode norms about gender, ability, and race in data visualizations and search interfaces. This critical scholarship complements projects with engineering colleagues at CSAIL to build more inclusive data representation systems. Drawing on work in feminist technoscience, disability studies, and the history and anthropology of computing, this dissertaiton weaves together different forms of HCI research to ask what work can or should be done by data representations across computational media.

Thesis Supervisor: Graham M. Jones
Title: Professor of Anthropology

Acknowledgements

In the process of completing this PhD, I have incurred so many debts that are impossible to repay. I am grateful first and foremost to my inimitable advisors, Graham M. Jones and Arvind Satyanarayan, for their unwavering support and endless guidance. They are the mentors, scholars, and humans I aspire to be in their relentless advocacy, rigorous research, and kindness. My committee members, Will Deringer and Eden Medina, have also played a huge supporting role in shaping the dissertation into what it is now — all mistakes are my own. I can only describe working with these faculty members as incredible in their willingness to help and work through things with me. Working through HASTS would have been impossible without the impeccable help of Karen Gardner, Kate Gormley, and Amberly Steward.

It is impossible to imagine what my graduate career would have looked like if not for the Visualization Group. They are as vibrant, funny, and sharp as they come: I am so grateful for their help fully cooking the many half-baked ideas I've had over the years. To Evan Peck, Aspen Hopkins, Angie Boggust, Dylan Wooton, Harini Suresh, and Wonyoung So: I am always better off after having a conversation with you all; each conversation is an embarrassment of riches. I hold special thanks to Alan Lundgard, Jonathan Zong, and Jamie Wong, who have been tireless co-authors and even better friends. Many of the projects described in these pages would have been impossible without you; being able to collaborate with you both has been one of the biggest privileges in my life. As friend, confidante, collaborator, and co-conspirator: you both have done so much more for me than you know.

I have been incredibly fortunate to have learned from so many wonderful undergraduate students at MIT and Wellesley, in particular Vesper Long, Yilia Qu, Tanya Yang, Kat Huang, and Mateo Monterde. Within HASTS and the Language & Technology Lab, I have had the good fortune of commiserating with and seeking advice about graduate school from fellow HASTSies Canay and Tom Özden-Schilling, Rodrigo Ochigame, Jamie Wong, Grace Kim-Butler, Marc Aidinoff, Clare Kim, Mariel García-Montes, Steven Gonzalez, Lauren Kapsalakis, Charlotte Minsky, Jia Hui Lee, Erik St. Gray, Beth Semel, Renée Blackburn, Gabrielle Robbins, and Lan Li. While disability studies is usually a niche area of work where it is almost impossible to find a community within the same institution, I have been blessed by my friendships and exciting exchanges with Di Wu and Timothy Loh. HASTS faculty like David Kaiser, Héctor Beltrán, Michael Fischer, Caley Horan, Robin Scheffler, and Dwai Banerjee have been crucial to my development as an STS scholar. Beyond HASTS, I have treasured the many stimulating conversations I've had with Lilly Chin, Catherine D'Ignazio, Milo Phillips-Brown, Marion Boulicault, Haley Schilling, Josh Pollock, and Geoffrey Litt. (Josh is single handedly responsible for my sanity in the last two days.) MIT is both an intellectually and emotionally taxing place, however: surviving it would have been impossible without the support of Sejal Patel, Leslie Langston, Bridget Burns, Kaitlyn Dwelle, Priya Moni, Leilani Gilpin, Emilia Simison, Annie Yuan Wei, Jane Hung, Parrish Bergquist, Shraddha Gupta, and Nicolette Bugher.

A fellowship at Harvard's Berkman Klein Center also led to some of my most cherished relationships in Cambridge, many of whom have helped shape formative ideas in this thesis about the social and legal dimensions of the internet. I thank Tom Zick, Kendra Albert, Mason Kortz, Andromeda Yelton, Kathy Pham, Sarah Newman, Jonas Kaiser, Sahar Massachi, John Maldonado, Boaz Sender, Akina Younge, Yves Daccord, Tunde Okunoye, Jad Esber, and the Ethical Tech Working Group. Mariel García-Montes and Rodrigo Ochigame in particular have been indispensable friends and guides at both MIT and Harvard. Outside of Cambridge, I have had the privilege of calling Kelly Wagman, Gracen Brilmyer, Levin Kim, Calvin Liang, Ranjodh Singh Dhaliwal, Alex Campolo, Meghan Kelly, Janice Kai Chen, and Emily Tseng both friends and collaborators. Ashley Shew has been indispensable to my development as a researcher working at the juncture of tech and disability, and I have found a warm and welcoming community with other disability studies scholars like Bess Williamson, Lydia X.Z. Brown, Damien Patrick Williams, Sandy Ho, Cynthia Bennett, and Joanne Woiak.

I began this journey into STS at Stanford, where I benefited from the advice of my faculty mentors Jessica Riskin, Paula Findlen, Adrian Daub, Londa Schiebinger, and Tom Mullaney. Hannah Leblanc, Alex Statman, Gina Russo Tam, and Anna Toledano were incredible office-mates and friends who are a pleasure to think with — our conversations have followed me beyond the beautiful Californian skies. I am also grateful to MIT Comparative Media Studies/Writing, particularly to Heather Hendershot, TL Taylor, Paul Roquet, and Eric Klopfer, who have engaged with my work very rigorously and compassionately. The feedback from a seminar at CMS/W has been indispensable as I write the last pages of this dissertation.

Weathering a pandemic would have been impossible without the support of Melissa Gardel, Kassidy Heath, Anja and Ian Calise, Thazin, Cynthia Lin, Chiara Giovanni, Mariel García-Montes, and Blakeley Payne. Vreny and Andreas Risi are the best in-laws anyone could ask for. To Yvonne: you are the light of my life and my biggest supporter. Professionally, you've coached me through every interview, every talk, every paper, and asked me the most incisive questions. As a sister, you help me feel less alone in the world — by reassuring me in times of darkness and by helping me enjoy the brightest moments. I love you. To Stephan: thank you for believing in me and in us. Reading through the chapters, I see your thinking and your generosity running through the pages. Life is sweeter walking with you and Laika together.

Finally, to my parents: 從我心底我要感謝爸媽在這三十年來的愛，支持和犧牲。若非如此，我就無法完成這篇博士論文。前人種樹，後人乘涼。

Funding for this dissertation was provided by MIT, the Programs for Digital Humanities, the Kelly Douglas Fund, the Schwarzman College of Computing, the Experiential Learning Opportunity (ELO) fund, the T.S. Lin Fellowship, and the Walter Rosenblith Presidential Fellowship. Outside of MIT, this research benefited from National Science Foundation Grant #1941577 and the Social Science Research Council's Social Data Dissertation Fellowship.

Contents

Introduction

There is no such thing as human-computer interaction, and this is a dissertation about it.

As a discipline, it is difficult to describe HCI synthetically by topic or by shared method. As it is reified in conferences like Computer-Human Interaction (CHI), as a subfield of computer science for hiring, and as a specialization legible to industrial researchers, the field has slowly taken on what George Stigler calls “an imperial science,” a term he used first to describe economics: “aggressive in addressing central problems in a considerable number of neighboring social disciplines, and without any invitation” (Stigler 1984). Talks within a single day at an HCI conference might employ methods from ethnographic observation to the design of machine learning systems, and the topics of inquiry can only be broadly described as computational or technological. Leading researchers come from a wide range of disciplinary backgrounds like anthropology, computer science, and cognitive psychology; the discipline also shares porous boundaries with other interdisciplinary fields such as science and technology studies.

Given the wide range of research contributions and scholarly norms between its constitutive scholarly communities, I rely here on Shaowen Bardzell’s (2010) definition of HCI as the “the design and evaluation of interactive systems,” which discusses both on critique-based (e.g., analyzing the unintended consequences of a design) and generative contributions (e.g., discovering solutions for new design problems). Even within a relatively simple and straightforward definition, HCI seems to study everything, everywhere, all at once: to chart and create the historical, present, and future development of sociotechnical systems. Given the diffuseness of the field, what epistemological and methodological approaches does it rely upon to become an academically legible discipline? How have proponents of the field been able to concretize it enough to be recognizable in computer science departments in the mid-20th century while remaining frustratingly diffuse?

While this introduction will only lightly touch upon the ill-defined field’s intellectual foundations to untangle how knowledge production occurs, each chapter in this dissertation showcases a different dimension of work that falls within this broad disciplinary remit. Like a gas, work in HCI morphs to fill the space it inhabits, which results in a large scope of scholarly investigation. These chapters take ethnographic observation, historical analysis, network analysis, interaction theory building, systems design, and user studies as tools to study the relationship between computational media and society. In so doing, this dissertation aims to showcase how these methods are mutually reinforcing rather than a product of the two cultures (Snow 2012): network analysis and classifying visualizations over half a million tweets can catch broad trends that would be impossible with close reading, and principles about accessible design drawn from qualitative research can be generative contributions towards designing more inclusive data visualization systems.

This integrated approach to qualitative research and the design of new computational systems is far from novel: in the late 1960s, ergonomics researchers like Brian Shackel (1969) outlined a clear agenda for the intersection between the human sciences and what was then termed man-computer interaction, an outgrowth of human factors research that itself came from the scientific management of labor processes (Williamson 2012, Grudin 2017). Interdisciplinary work of this flavor has been long pioneered by scholars of computation like anthropologist Lucy Suchman (1985) and computer scientist Paul Dourish (2001). This genre of contribution has been professionalized for industrial and academic researchers at conference venues such as Computer-Supported Cooperative Work (CSCW). Today, this blend of social research on computational ecosystems continues to thrive at industry research labs (e.g., Microsoft's Social Media Collective) and schools of information and computing. A foundational paper that many computer scientists and historians of technology point to when it comes to the historical genesis of HCI is J.C.R. Licklider's "Man-Computer Symbiosis" (1960), which proposes a "cooperative interaction between men and electronic computers," where users set up the goals, hypotheses, and evaluation processes for a particular problem while computers are able to complete automated tasks at a speed and volume that would be impossible for humans.

The goal of this endeavor is ultimately to optimize for human performance and efficiency: to complete tasks quickly and advance human productivity, which has precursors in early 20th century American history. Proponents of scientific management like Frederick W. Taylor and Frank and Lillian Gilbreth, for example, conducted time and motion studies in order to understand how human operators interacted with industrial equipment. Taylor himself suggested screening workers' bodies for size and fitness with the machine they would work with (Williamson 2012), while the Gilbreths pioneered micro motion studies as a way of understanding and regulating worker productivity (Nadworny 1957, Gilbreth and Gilbreth 1917). These studies used the chronocyclegraph (Figure 1) to chart the workers' movements, which is a small electrical lamp attached to a worker's fingers which is then photographed with a multiple exposure method in order to chart the worker's movements (Figure 2). Through repetition, the marks that the subjects made on the photograph became brighter, and the gridded backdrop made it possible for the Gilbreths to calculate the precise distance between their movements. The images would then be transposed into 3D wire sculptures, which could show a subject's progress as they learned how to complete a task (Figure 3).

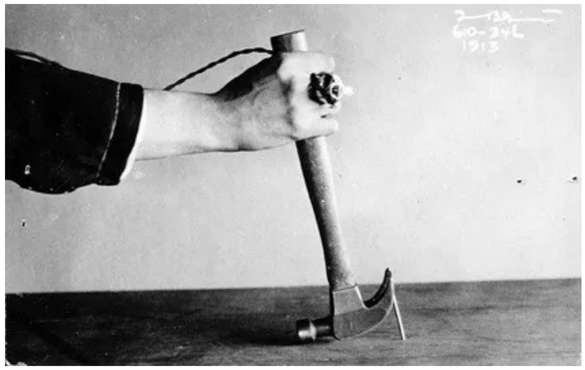

Figure 1. “Method of attaching light to hand for cyclograph pictures. Light ring on hand holding hammer while pulling a nail. 1913.” Frank B. Gilbreth Motion Study Photographs (1913-1917). The Kheel Center for Labor-Management Documentation and Archives.

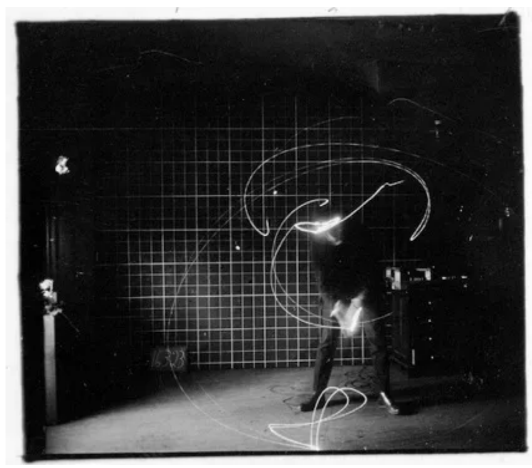

Figure 2. “Chronocyclegraph of golf champion: Francis [?] , circa 1915.” Frank B. Gilbreth Motion Study Photographs (1913-1917).

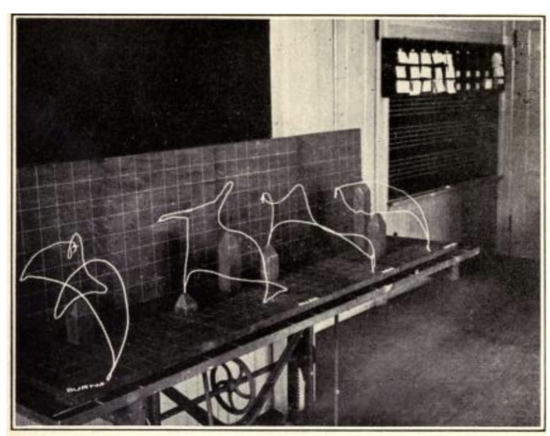

Figure 3. “First photograph of wire models showing one man’s progress of learning paths of least waste.” (Gilbreths 1917, 92).

While most industrial consultants never seriously considered disability as a category in the study of industrial work, the Gilbreths' studies of World War I veterans in rehabilitation hospitals (Brown 2002, cited in Williamson 2012, 215) later paved the way for "ergonomic" designs that accounted for "human factors" (Mead and Wulfeck 1952). This attention to human variation was in part guided by the problem of designing military equipment for a large population of American soldiers in the 1940s, where anthropometric methods created by interdisciplinary teams of doctors, psychologists, anthropologists, and engineers – to create a "standard man" as a baseline for all variations of an object's design (Dreyfuss 1967). These universal benchmarks could then be mobilized to design planes and submarines with optimally placed levers, buttons, and doors (Williamson 2012).

As soldiers returned from WWII, however, ensuring that the design of everyday objects could be accessible to returning veterans quickly became "allied to a notion of national efficiency and post-war reconstruction" (Hayward 1998); the goal was to seamlessly integrate veterans into a new, booming economy as civilian workers. Soldiers who returned with significant injuries underwent medical rehabilitation aimed at normalizing their bodily functions so that they could become both productive workers and citizens (Hamraie 2016). In the words of design historian and ethnographer Aimi Hamraie, "given the right equipment and appropriate built environment, injured soldiers could work in factories, drive automobiles, and realize the American Dream. Disabled people, in other words, could become normal by design" (ibid. 2016, 6). The ultimate goal of these design and rehabilitation projects was to optimize the creation and maintenance of a large, functioning workforce that could span professions and ability. That a computer could be as easily used by an artist as it could be by a programmer to complete tasks with augmented abilities was one manifestation of this idea (e.g., Sutherland 1963); computers could suffuse almost every form of work in ways that would lessen the repetitive burdens on the worker and augment their ability to complete complex tasks.

In Douglas Engelbart's 1962 report on augmenting human intellect to the Air Force's Director of Information Sciences (inventor of the computer mouse and putative founder of HCI), he argues that "improving intellectual effectiveness" should be a prominent target of government research investment: "after all, we spend great sums for disciplines aimed at understanding and harnessing nuclear power. Why not consider developing a discipline aimed at understanding and harnessing 'neural power?' In the long run, the power of the human intellect is really much the more important of the two" (Engelbart 1962, 132). The optimization of intuitive, easy to use systems was a core concern for early human factors research, which is a legacy that has extended towards HCI (e.g., "to lighten cognitive load"); this focus on intuitive or easy to use technologies relies upon notions of techno-solutionism and universalism that this dissertation tries to unsettle. In other words, the case studies within this dissertation use COVID-related visualizations, memes, and accessible visualization design in order to show how research in HCI is historically contingent, socially situated, and materially embodied.

Instead of asking how to make systems more intuitive or to optimize them for speed, the case studies in this dissertation instead interrogate *how* or *why* an interface becomes "intuitive" (and to whom). The Gilbreths' motion studies, like the eye-tracking and quantitative user studies of modern HCI research, are not objective renderings of how a person interacts with technology; these studies are themselves a product of historical conditions and embodied experience (Daston and Galison 2007). Technology that is "better designed for everyone" will not solve social problems, and the totalizing sweep of universal design minimizes (and often erases) the problems shouldered by multiply marginalized people, as one-size-fits-all solutions often only reflect the experience of the most privileged users. The military and political origins of computing and the design of computational interfaces too suffuse the field's values and goals; research on computer-supported cooperative work has historically optimized for productivity and efficiency in ways that often overlook the social ecosystems in which these technologies live.

HCI and design research need not position technical systems as universally applicable solutions to social problems; scholarship and practice in this field can be (and has been) informed by critical engagements with technology and the social conditions in which they are developed. Work drawn from cognitive psychology and neuroscience in HCI may break down a person's individual engagement with a system in more biological terms, but work from history and anthropology contextualizes how certain technologies work in communities of people with overlapping (or conflicting!) interests and goals. As a participant observer and an observant participant within an HCI lab over the last three years, I have drawn liberally from my conversations with fellow lab mates to interrogate the mythical value of "universal access" within data visualization and HCI, and to show how data representations are personally and culturally situated. These qualities can often make these data representations and technical systems into objects of controversy and critique that is ultimately about power. Drawing on work in feminist technoscience, disability studies, and the history and anthropology of computing, this work attempts to weave together different forms of HCI research to ask what work can or should be done by data representations across computational media.

Why write about human-computer interaction?

Historians of computing have taken a sweeping approach to understand the rise of tech and the academic-military-industrial complex. Tech metropolises like Palo Alto are hardly made up of bootstrap operations funded and powered by lone wolf coders; historians like Margaret O'Mara (2004) and Paul Edwards (1997) have shown how university tech behemoths like MIT and their industrial spin offs were the products of massive investment from Cold War defense budgets. These planned communities — from talent recruitment to material infrastructure — were shaped both by national defense priorities and postwar economic planning. As Edwards describes in his history of computing and scientific subjectivity titled *The Closed World,* "we can make sense of the history of computers as tools only when we simultaneously grasp their history as metaphors in Cold War science, politics, and culture" (Edwards 1997, ix). Edwards argues that the government's exorbitant investment in computing technologies spoke less to the computers' actual functions and more to how the technologies could ideologically (and aspirationally) provide a sense of safety and control to military actors. Computational machines promised a future filled with unified systems of command that could eliminate human error.

Scientists and Cold War strategists created this future via the "closed world" — one where the mess of real world problems could be sanitized and ordered into bounded and rational space. This illusion of control centered around "problems of *human-machine integration:* building weapons, systems, and strategies whose human and machine components could function as a seamless web…[where] computers would automate or assist tasks of perception, reasoning, and control in integrated systems" (Edwards 1997, 1). This approach to studying and developing systems (human and technological) as self-regulating machines was the intellectual cornerstone of many allied and overlapping disciplines, like cognitive psychology and cybernetics. Within this world, computational technology could function as prostheses for the human sensorium, where cyberneticians like Norbert Wiener at MIT explored modes of sensory substitution through objects like the hearing glove — a device that would translate speech extracts through pressure on one's fingers (Figure 4; see also Mills 2011). The production of tactile sensation as a way of making language and computers accessible to people with disabilities is an important development in human-computer interaction that reveals the phenomenal and interactional consequences of language, signal processing, and tactility as a method of conveying information. Historian Orit Halpern's account of cybernetic reformulations of perception also shows an underlying tension between "envisioning sensory infrastructure as subjective and mimetically prosthetic or networked and autonomous" (Halpern 2012, 234). Computational tools could either be a way of subjectively enhancing the human sensorium *or* be part of an objectively ordered system of people and machines that could seamlessly transfer information.

The cybernetic narrative that technologies could remediate or augment the human senses at MIT was reinforced following an injury that left Wiener with a broken hip in 1962. Following this injury, doctors and researchers at the Massachusetts General Hospital and Harvard Medical School worked with Robert Mann (an MIT professor of mechanical engineering) to construct a small power system that could be used to send electrical signals from the brain to an artificial arm (MIT Museum 2011; Figure 5). The power system itself was based on an earlier design for a lightweight system that Mann developed in the 1950s for the Sparrow and Hawk missiles (Wright 2006). In some ways, Mann's 40 year career at MIT is the perfect embodiment of the transition of human factors research from military applications to remediating the effects of war: in a public lecture at MIT, he recounts his career switch from "powering rockets to powering people" (Mann 1997) — an ethos he brought to the design of the flagship course for MIT mechanical engineers, 2.007 (Design and Manufacturing I).

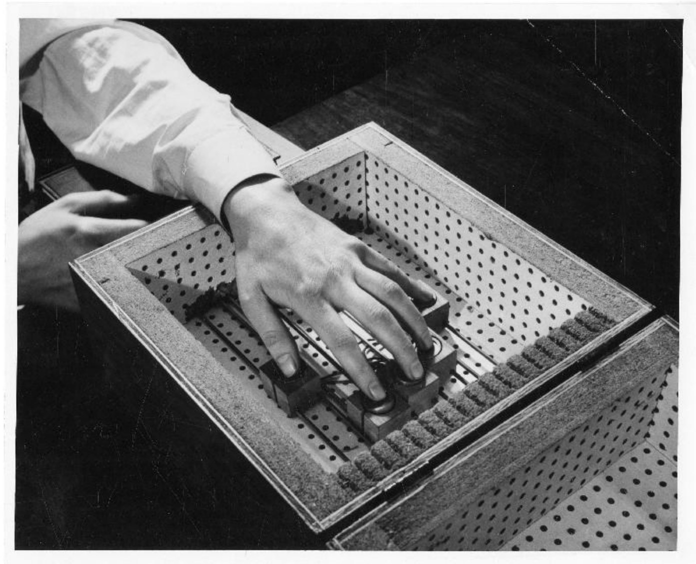

Figure 4. Prototype for the Hearing Glove.
Photo by Alfred Eisenstaedt, May 24, 1949. Norbert Wiener papers, MIT Museum.

Figure 5. Man in a suit and tie using an early prototype of the Boston Arm to pick up the telephone. Collection on the Boston Arm (Boston Elbow), MIT Museum.

The idea that using and manipulating a programmable machine is a sensuous experience, however, has a longer history. In 1844, Augustus de Morgan — the founder of the London Mathematical Society and tutor to Ada Lovelace — wrote to his student's mother praising Lady Lovelace's mathematical prowess but cautioned that her aptitude and work would be too onerous for her body (Mayne 1929). In a 1984 essay in *Scientific American*, computer scientist Alan Kay — the creator of the Dynabook, the conceptual predecessor to the laptop computer — likened computers to musical instruments that, drawing on Da Vinci, "shape the invisible" (Kay 1984, 53). Unlike clay that can be shaped simply by "shoving both hands into the stuff," Kay argued that using a computer as a programmer is as "detached from human experience as a radioactive ingot being manipulated remotely with buttons"; crossing this gap requires attention to the design of user interfaces, as "what is presented to one's senses *is* one's computer" (ibid., 54). This feeling of direct sensory access is encapsulated in later studies by cognitive scientists Edwin Hutchins, James Hollan, and Donald Norman (1985) as "direct manipulation," which describes the "feeling of directness [...] between the user's intentions and the facilities provided by the machine," which requires the "system provide representations of objects that behave as if they are the objects themselves." The ability to conjure and enhance the "feeling of directness" has been a cornerstone in tech development for personal and handheld devices in recent years, encapsulated best by the print slogan for the first iPhone: "touching is believing" (Figure 6).

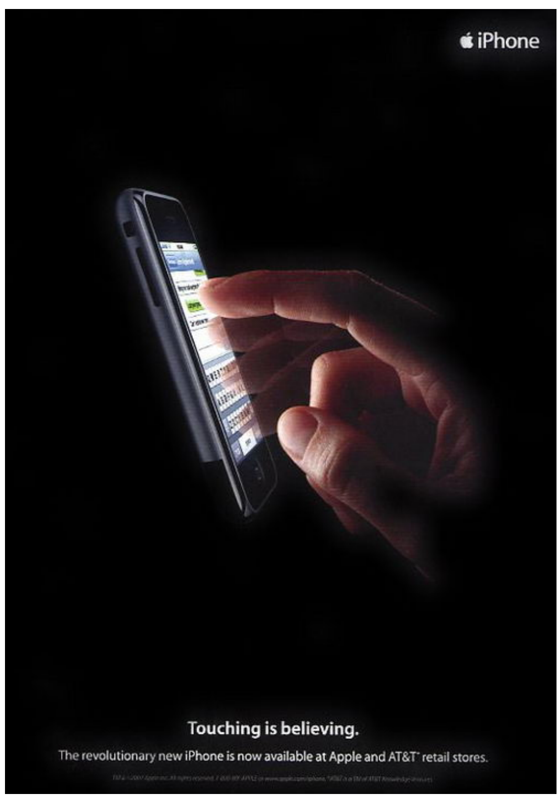

Figure 6. Print ad for the first iPhone in *Details Magazine,* August 2007.

Such are the varied foundations on which human-computer interaction research is built in the United States. From the development of ergonomic devices to measuring human perception to optimize computational systems, a great deal of this origin story has thus far been told primarily through prominent figures in Cold War science like Vannevar Bush and Douglas Engelbart. Their influence, and that of Stanford and MIT, is indisputably significant; the computer mouse, multiplayer video games, and touch screens are important milestones in HCI research that were developed at or incubated by these universities. However, these are also the glorious stories that departments recount to themselves to cement their already outsized intellectual reputations. They are the same accounts that become inculcated in students in this lineage who also pursue research within these fields; nearly every course on human-computer interaction or design at these universities open with Vannevar Bush's 1945 essay *As We May Think* (which described the memex, or an early hypertext system) within introductory lectures or assign it as reading. How did this canon of putative founders of HCI come about, and what are alternative ways of envisioning the intellectual origins of the discipline? One way of understanding the history of HCI and its intellectual foundations is through connecting two bodies of interpretive inquiry: one in design history and computational media (e.g., Williamson 2012, Hamraie 2016) and the history of computing and labor (e.g., Abbate 2000, 2012; Ensmenger 2012; Hicks 2017; Gray and Suri 2019). Refracting issues of interaction and design through questions about labor and the future of work allows me to productively draw upon extant historical sources and contextualize them within a broader political economy of data and computing.

My research will depart from conventional histories about computing by historicizing and culturally relativizing claims about the universality and immediacy of touch. Previous histories of HCI have tended to peg changes in interaction theory with the introduction of new technologies like the mouse or the graphical user interface (Dourish 2001), and much of this work focuses on the underlying phenomena — like cognitive and learned characteristics — that "give rise to the feeling of directness" (Hutchins, Hollan, and Norman 1985; Frohlich 1993). "Directness"— that is, the distance between the user's expressed intention and their ability to actually execute this action within a computing system. Anthropologists and historians have recently shown how this principle reflects tacit knowledge about how computers ought to work (Collins 2001; Petrick 2015), and while early HCI literature heavily emphasizes language around how the body interacts with a computer interface, I hope to build on and contribute to the lively, growing body of work on the history of touch and haptic interfaces (Maines 2001; Paterson 2007, 2016; Plotnick 2012, 2018; Galloway 2017; Scherffig 2018; Velasco and Obrist 2020).

Scholars in sensory studies have turned their investigations towards understanding how sensory experiences — the scent of a perfume, the intimacy of touch — are themselves fields of cultural elaboration (Corbin 1986; Stoller 1997; Howes 2003; Classen 2012; Pink 2015). These anthropologists posit a culturally specific narrative of the sensorium against a purely biological, neuroscientific one that maps human senses to specific parts of the brain or the body. Similar work by philosophers articulate how different ways of understanding human perception impacts our ideas about what is knowable (Merleau-Ponty 2013; Stokes, Matthen, and Biggs 2014; Barwich and Chang 2015). In particular, historians and sociologists have written extensively on the Molyneux problem — a thought experiment made famous by John Locke about the relationship between tactile and visual knowledge — to understand the place of blindness in philosophy as a counter to the centrality of vision in scientific practice (Riskin 2002; Jütte 2004; Paterson 2016). This dissertation builds on literature in sensory studies to show how accounting for disability in the design of tactile objects requires a renewed attention to non-visual modalities and their cultural histories (Harris 2007; Ingold 2011; Paterson and Dodge 2012). By bringing sensory studies to the history of computing, I hope to show how the design of tactile objects embeds within it narratives about the senses as a culturally and technologically mediated experience (Howes 2005; C. A. Jones 2006).

On critical technical practice

While I acknowledge the influence that MIT and its affiliates have had on the history of computing and HCI, buying into this narrative uncritically often reproduces social inequality and erases the hidden labor that sustains — or liberatory alternatives to — our existing tech ecosystems. As Daniela Rosner (2018) writes in her critical history of design, elevating long-silenced narratives about design and their methodologies helps technologists challenge existing paradigms to imagine more equitable worlds. Similarly, in their history of Latin American informatics, Rodrigo Ochigame (2020) shows how investigating historical experiments with different approaches to information sharing (in their case, with Cuban socialism and Latin American liberation theology) can help us "imagine new ways to organize information that threaten the capitalist status quo — above all, by facilitating the wide circulation of the ideas of the oppressed." Prison abolitionists and anti-capitalist scholars have also described how utopian visions, critical histories, and pragmatic reforms all "inform one another in order to produce utopian destinations that have accessible waystations" (Green 2021). Critical engagements with speculative futures and counter-histories — particularly those that challenge extant narratives about *how* the status quo came to be — are important for proactively challenging how computational technologies reinforce, rather than resist, structures of power. Feminist STS scholars Frances Bronet and Linda Layne suggest that scholars in the humanities and social sciences should not "sit back and offer post facto critiques of new technologies, but [instead] … intervene to proactively influence design" (Bronet and Layne 2010, 179).

The relationship between critical, historical work and the development of new computing systems is a fundamentally dialectical one. In The *Mangle of Practice*, STS scholar Andrew Pickering (1995) recounts the story of a physicist, Donald Glaser, who was trying to construct a new instrument to detect so-called "strange particles" in the 1950s. While the machine Glaser initially developed did not register any particle tracks, he constantly reconfigured both his machine and the parameters he was measuring, which slowly began to garner new results. For every sequence of new obstacles — where the experiment would return little to no particle traces — Glaser would metaphorically readjust the lens, change another parameter, and perform the experiment again. This dialectical pattern of resistance followed by accommodation — where a scientist adjusts the lens of a microscope, only for the object of study to blur the image, thus prompting readjustment — is what Pickering calls the mangle of practice. Science as a knowledge enterprise is generated by this mangling of human and machine, a metaphor which accentuates both the instrumental and messy aspects of science. The constant form of resistance and readjustment is an interesting thread to follow throughout the history of interfaces, interaction, and gesture, as HCI research provides an ideal case study for investigating the mangle of practice in computing.

Methods and collaboration

The studies in this dissertation use methods across the humanities and social sciences to study social media and accessible data representations. This is partially out of necessity — the pandemic made in-person archival and field research impossible — but the methods that we ultimately used in each study were first and foremost a reflection of what suited the data best rather than an exercise in demonstrating methodological competence. Additionally, with the exception of the introduction and conclusion, each chapter within this dissertation is a co-authored piece. I have had the good fortune of finding great collaborators on these projects who have become even greater friends. In particular, my exploration of computer science was only possible with the encouragement of Alan Lundgard and the generosity of Arvind Satyanarayan; the ability to truly collaborate across disciplines has been one of the greatest privileges of my time at MIT. This brings me to answer a broader question about my own role within these collaborations, and to clarify what my and my collaborators' contributions were to each chapter. Even these can be blurry descriptions, as the projects were successful because many of these contributions were mutually reinforcing; the final projects were greater than the sum of their individual contributions. I have integrated notes on the division of labor for each of these projects in the section below.

Dissertation structure

This dissertation consists of two majors sections: one on social media, and one on accessible data representations. In the first section on social media, I present two case studies about the dissemination of COVID-related data across the internet in the US (Facebook and Twitter) and in China (Weibo). The contributions here are primarily descriptive and comparative (Beaudouin-Lafon 2000): they analyze discourse across these platforms about COVID over time and the social dynamics that make these ideas trenchant, but they do not suggest immediate changes to the design of new technologies. To be clear, however, when I call research descriptive, it does not mean that it is a passive accounting of events that are then subject to some form of analysis. Data collection in qualitative research is itself an active, iterative process where the researcher simultaneously collects and produces data, as the researcher's presence and perspective indelibly shapes how a dataset is structured and analyzed. Semi-structured interviews all require some degree of extemporaneous speaking and its focus changes circumstantially, and analyzing these data requires personal contextualization that often only the ethnographer can provide. In the words of sociologists Jessica Calarco and Mario Small: "given that the ethnographer writes every one of the hundreds of thousands of words that constitute a typical set of field notes, the observer is inescapably embedded in the data themselves" (Small and Calarco 2022, 12).

This form of limited, purposive sampling can often be epistemologically unsatisfying; to respond to these claims, qualitative researchers might feel pressured to incorporate features consistent with quantitative research that include things identifying "representative" neighborhoods or selecting respondents "at random." Mario Small (2009) argues that this corrective actually makes the insights from qualitative research less potent. Instead of fixating on the data as being representative or generalizable, qualitative study design should instead be focused on *saturation* (where each additional case reveals less and less surprising information) and on *juxtaposition* (evaluating whether or not new insights are generally consistent with existing scholarship). The point is not to replicate qualitative studies asking the same questions in order to establish the validity of the results, but to evaluate how a particular study is consistent with (or departs from) existing scholarship. The case studies below are hardly discontinuities in research on social media and data visualizations, but the data we collected and analyzed bring up new questions about the public rhetoric of science and the design of visual information.

The first case, **Let's Go, Baby Forklift!**, describes how the Chinese government co-opted a celebrity fandom forming around anthropomorphized construction vehicles deployed to address the COVID crisis in Wuhan. At the start of the pandemic, residents posted live blogs and videos showing the city's hospitals being overwhelmed, and some of the more devastating posts detail users slowly succumbing to the disease until their accounts run cold. In response to the crisis and burgeoning criticisms circulating widely on Weibo and WeChat, the Chinese Communist Party mobilized to place nearly 50 million people in Hubei province in quarantine and began constructing two emergency hospitals that would be operational in just over a week (Qin 2020). The construction of these hospitals was subject to a round-the-clock livestream where users could "witness" (见证) the entire process, after which some of the commenters began referring themselves as "supervisors" (监工) ready to take over the last shift. In three days, the livestream had received over 190 million views and 77,000 comments. A curious fandom emerges from this enormously popular livestream: commenters begin anthropomorphizing the construction vehicles — most notably, a yellow forklift — and use them to create idols that sometimes displaced singer/actor idols on social media popularity charts.

Drawing on methods in cultural and linguistic anthropology, this study provides an account of how the memeification of these construction vehicles and their fandoms become leveraged by Chinese state social media to consolidate state power. I began this project with the intention of examining the circulation of visual information and memes about the pandemic in China working with Vesper Long. Throughout the summer, Vesper and I worked together to iteratively analyze and collect data; she harvested hundreds of screenshots of memes and reactions that Graham Jones, Vesper, and I discussed with Chinese social media experts Jamie Wong and Di Wu. My primary role in this project was to facilitate and guide data collection, and to draw out themes that I would later discuss and refine with the rest of the team. The analysis for this project was truly conducted collaboratively across multiple time zones; over the course of the project, there were hundreds of pages of memos and screenshots that we ultimately distilled into an article.

In the second case, **Viral Visualizations**, the study follows the dissemination of COVID-related data visualizations and the social media discussions about how the pandemic is unfolding. The rhetorical force of these data representations has been apparent especially at the start of the pandemic, where visualizations like “Flatten the Curve” and interactive dashboards abounded alongside 24/7 reporting. Given the ubiquity of these visualizations, how did users across Facebook and Twitter rhetorically deploy visualizations to frame public arguments about public health, economic recovery, and the role of data in policy making? How can data and their visual representations be manipulated in online discussions? Ultimately, this case study shows that data visualizations are not simply tools by which people understand the epidemiological events around them; they instead mediate relationships between public health officials, epidemiologists, the news media, and the broader public.

Drawing on sociological research on news media in conservative Christian communities (Tripodi 2018) and Tea Party supporters in Louisiana (Hochschild 2016), this chapter unpacks the data practices of coronavirus skeptics to argue that injunctions towards increased data literacy may not succeed, as the tools and frameworks that “data literacy” might entail have been weaponized by groups seeking to cast doubt on official COVID datasets. These skeptics employ many of the same tropes that might be used in scientific papers or data journalism pieces; they implore readers to question the circumstances around which data are collected and the structure of how it is then presented. Visualizations — especially as created for public health officials or news media — are often misleading; better to return to the most basic forms of data like tables (or better yet, to collect the data yourself). As an STS scholar, the most fascinating artifact of this study is the extent to which these coronavirus skeptic communities similarly deploy traditional STS arguments: citing Karl Popper, every form of scientific conclusion should be met with continuous falsification in order to be properly rigorous; citing Thomas Kuhn, there are sufficient data that undermines current understanding of a scientific phenomenon that it is time for a new paradigm shift.

Methodologically, this study employs both qualitative data collection of social media posts (depth) alongside computational analysis of half a million tweets and over 400,000 data visualizations (scale) in order to create a more synthetic view of COVID data-related discourse on Facebook and Twitter. In this study, I collected and analyzed all of the qualitative data from Facebook with Gabrielle Inchoco and Graham Jones. Tanya Yang and Arvind Satyanarayan collected the Twitter data and created the network visualization. They also compiled the image dataset and analysis for the chart classification schema; I contextualized and analyzed the results.

The second half of this dissertation is focused on the design of accessible data representations. In the third chapter, **Sociotechnical Considerations**, we introduce a set of principles on how best to approach the design of these systems. Drawing on work in critical disability studies and STS, we argue that researchers should be both critical of technological interventions and clear about how the research contributes to disabled communities given the history of extractive scholarship. Aside from the basics of fairly compensating disabled collaborators and adhering to extant accessibility guidelines, our paper also introduces ways of thinking about technological resolution and access to understand how tactile visualizations can effectively encode information. We use these principles to examine a small case study we conducted through a joint workshop at MIT and the Perkins School for the Blind. In this project, Alan Lundgard and I both attended the workshop and collected qualitative data from those encounters, and through our ensuing discussions, iteratively distilled the data we had collected into the major themes.

In the fourth chapter, **Rich Screenreader Interactions,** we present a design space for creating interactive data representations for screen readers with new prototypes. Current best practices for accessible data visualization on the web usually come in two forms: captions that summarize the data, or raw data tables where screen reader users can tab through each cell individually. This paper presents an alternative: an interactive system where screen reader users can zoom in and out of visualizations — in other words, from a high level summary of trends down to individual data points — and compare trends across multiple graphs. Users can also selectively attend to their points of interest by jumping directly into a particular slice of the data and move around directionally (e.g., in a graph that maps a person's age to height, one could directly jump to all people who were above 6 feet tall and processually move through their ages). In sharp contrast to data tables (even with filtering and sorting options), our interactive system provides screen reader users with the ability to move dynamically across a graph. This study was constructed with the principles from "Sociotechnical Considerations" in mind. I participated in and took field notes on the iterative design process between Daniel Hajas (our blind collaborator at University College London) and our team (Jonathan, Alan, Arvind, Joon, and myself). I also co-ran the user studies with Jonathan. Jonathan engineered the prototypes with some assistance from Alan and Joon. The design dimensions for the paper arose from discussions between Jonathan, Alan, Arvind, and myself.

The final two papers represent a first attempt at leveraging qualitative work on disability and the senses into the design of new data representation systems. Taken together, the four chapters refract different ways of approaching problems in human-computer interaction that reflect its tangled disciplinary history. I hope to extend the work from this dissertation into a more synthetic project that investigates what I call the *life-cycle of data representations* — the process by which datasets are curated, cleaned, visualized, edited, circulated, and manipulated — to understand the persuasive force of visual information in multimodal media. The historical work on HCI and its theoretical foundations will inform my collaborative research with computer scientists about multisensory grammars for data representations. By combining the history of design and disability and these multi-year collaborations, I hope to expand these chapters into a book project that studies visualization as a community of practice, as a media ecology, and as a form of multisensory knowledge. Exploring haptic media like tactile maps challenges the ocularcentrism of conventional visualization practices, and it is an opportunity for media scholars and critical makers to productively interrogate the relationship between the senses and interface design.

References for Introduction

Abbate, Janet. *Inventing the Internet*. Cambridge, MA: MIT Press, 2000.

———. *Recoding Gender: Women's Changing Participation in Computing*. Cambridge, MA: MIT Press, 2012.

Bardzell, Shaowen. "Feminist HCI: Taking Stock and Outlining an Agenda for Design." In *Proceedings of the 28th International Conference on Human Factors in Computing Systems - CHI '10*, 1301. Atlanta, Georgia, USA: ACM Press, 2010. https://doi.org/10.1145/1753326.1753521.

Barwich, Ann-Sophie, and Hasok Chang. "Sensory Measurements: Coordination and Standardization." *Biological Theory* 10, no. 3 (September 1, 2015): 200–211. https://doi.org/10.1007/s13752-015-0222-2.

Beaudouin-Lafon, Michel. "Instrumental Interaction: An Interaction Model for Designing Post-WIMP User Interfaces." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 446–53. CHI '00. New York, NY, USA: ACM, 2000. https://doi.org/10.1145/332040.332473.

Bronet, Frances, and Linda L. Layne. "Teaching Feminist Technology Design." In *Feminist Technology*, edited by Sharra Louise Vostral, Kate Boyer, and Linda L. Layne. University of Illinois Press, 2010.

Brown, Elspeth. "The Prosthetics of Management." In *Artificial Parts, Practical Lives: Modern Histories of Prosthetics*, edited by Katherine Ott, David Serlin, and Stephen Mihm, 249–81. New York: NYU Press, 2002.

Classen, Constance. *The Deepest Sense: A Cultural History of Touch*. Studies in Sensory History. Urbana, IL: University of Illinois Press, 2012. http://public.eblib.com/choice/publicfullrecord.aspx?p=3414197.

Collins, Harry M. "What Is Tacit Knowledge?" In *The Practice Turn in Contemporary Theory*, edited by Theodore R. Schatzki, Karin Knorr-Cetina, and Eike von Savigny, 107–19. London: Routledge, 2001.

Corbin, Alain. *The Foul and the Fragrant: Odor and the French Social Imagination*. Cambridge, MA: Harvard University Press, 1986.

Daston, Lorraine, and Peter Galison. *Objectivity*. Cambridge, MA: Zone Books, 2007.

Dourish, Paul. "History of Interaction." In *Where the Action Is: The Foundations of Embodied Interaction*. Cambridge, MA: MIT Press, 2001.

Dreyfuss, Henry. *The Measure of Man: Human Factors in Design*. 2nd ed. New York: Whitney Library of Design, 1967.

Edwards, Paul N. *The Closed World: Computers and the Politics of Discourse in Cold War America*. Reprint edition. Cambridge, MA: MIT Press, 1997.

Engelbart, Douglas C. "Augmenting Human Intellect: A Conceptual Framework." Summary Report. SRI Project Number 3578. Washington, D.C.: Director of Information Sciences, Air Force Office of Scientific Research, 1962.

Ensmenger, Nathan L. *The Computer Boys Take Over: Computers, Programmers, and the Politics of Technical Expertise*. Computers, Programmers, and the Politics of Technical Expertise. Cambridge, MA: MIT Press, 2012.

Frohlich, David M. "The History and Future of Direct Manipulation." *Behaviour & Information Technology* 12, no. 6 (November 1, 1993): 315–29. https://doi.org/10.1080/01449299308924396.

Gaboury, Jacob. *Image Objects: An Archaeology of Computer Graphics*. Cambridge, MA, USA: MIT Press, 2021.

Galloway, Alexander R. *The Interface Effect*. Cambridge, UK; Malden, MA: Polity, 2017.

Gilbreth, Frank Bunker, and Lillian Moller Gilbreth. *Applied Motion Study; a Collection of Papers on the Efficient Method to Industrial Preparedness*. New York, Sturgis & Walton company, 1917. http://archive.org/details/appliedmotionstu00gilbrich.

Gray, Mary L, and Siddharth Suri. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Boston, MA: Houghton Mifflin Harcourt, 2019.

Grudin, Jonathan. *From Tool to Partner: The Evolution of Human-Computer Interaction*. Morgan & Claypool, 2017. http://www.morganclaypool.com/doi/10.2200/S00745ED1V01Y201612HCI035.

Halpern, Orit. "Cybernetic Sense." *Interdisciplinary Science Reviews* 37, no. 3 (September 2012): 218–36. https://doi.org/10.1179/0308018812Z.00000000018.

Hamraie, Aimi. "Universal Design and the Problem of 'Post-Disability' Ideology." *Design and Culture* 8, no. 3 (September 2016): 285–309. https://doi.org/10.1080/17547075.2016.1218714.

Harris, Mark, ed. *Ways of Knowing: Anthropological Approaches to Crafting Experience and Knowledge*. New York; Oxford: Berghahn Books, 2007.

Hayward, Stephen. "'Good Design Is Largely a Matter of Common Sense': Questioning the Meaning and Ownership of a Twentieth-Century Orthodoxy." *Journal of Design History* 11, no. 3 (January 1, 1998): 217–33. https://doi.org/10.1093/jdh/11.3.217.

Hicks, Marie. *Programmed Inequality: How Britain Discarded Women Technologists and Lost Its Edge in Computing*. Cambridge, MA: MIT Press, 2017.

Hochschild, Arlie Russell. *Strangers in Their Own Land: Anger and Mourning on the American Right*. New York: New Press, 2016.

Howes, David, ed. *Empire of the Senses: The Sensual Culture Reader*. Sensory Formation Series. Oxford; New York: Berg, 2005.

———. *Sensual Relations: Engaging the Senses in Culture and Social Theory*. Ann Arbor, MI: University of Michigan Press, 2003.

Hutchins, Edwin L, James D Hollan, and Donald A Norman. "Direct Manipulation Interfaces." *Human-Computer Interaction* 1 (1985): 311–38.

Ingold, Tim. "Worlds of Sense and Sensing the World: A Response to Sarah Pink and David Howes." *Social Anthropology* 19, no. 3 (2011): 313–17.

Jones, Caroline A., ed. *Sensorium: Embodied Experience, Technology, and Contemporary Art*. Cambridge, Mass: MIT Press, 2006.

Jütte, Robert. *A History of the Senses: From Antiquity to Cyberspace*. Translated by James Lynn. 1 edition. Cambridge, UK ; Malden, MA, USA: Polity, 2004.

Kay, Alan. "Computer Software." *Scientific American* 251, no. 3 (1984): 52–59.

———. "User Interface: A Personal View (1989)." In *Multimedia: From Wagner to Virtual Reality*, edited by Randall Packer and Ken Jordan, 1st ed., 121–31. New York: Norton, 2001.

Licklider, J.C.R. "Man-Computer Symbiosis." *IEEE Transactions on Human Factors in Electronics*, March 1960, 4–11.

Maines, Rachel P. *The Technology of Orgasm: "Hysteria," the Vibrator, and Women's Sexual Satisfaction*. Revised ed. edition. Baltimore; London: Johns Hopkins University Press, 2001.

Mann, Robert. "A Half-Century Portfolio on Engineering Design." 1997 Distinguished Alumni/ae Lecture Series, MIT Department of mechanical Engineering, March 12, 1997. https://news.mit.edu/1997/meche-0305.

Mattern, Shannon. “Mission Control: A History of the Urban Dashboard.” *Places Journal*, March 9, 2015. https://doi.org/10.22269/150309.

Mayne, Ethel Colburn. *The Life and Letters of Anne Isabella, Lady Noel Byron from Unpublished Papers in the Possession of the Late Ralph, Earl of Lovelace*. London, Constable & Co., 1929.

Mead, Leonard C., and Joseph W. Wulfeck. “Human Engineering: The Study of the Human Factor in Machine Design.” *The Scientific Monthly* 75, no. 6 (1952): 372–79.

Merleau-Ponty, Maurice. *Phenomenology of Perception*. Translated by Donald Landes and Taylor Carman. 1 edition. Abingdon, Oxon ; New York: Routledge, 2013.

Mills, M. “On Disability and Cybernetics: Helen Keller, Norbert Wiener, and the Hearing Glove.” *Differences* 22, no. 2–3 (January 1, 2011): 74–111. https://doi.org/10.1215/10407391-1428852.

MIT Museum. “Subjects: Boston Arm (Boston Elbow),” 2011. https://webmuseum.mit.edu/detail.php?module=subjects&type=related&kv=1167.

Nadworny, Milton J. “Frederick Taylor and Frank Gilbreth: Competition in Scientific Management.” *Business History Review* 31, no. 1 (ed 1957): 23–34. https://doi.org/10.2307/3111727.

Ochigame, Rodrigo. “Informatics of the Oppressed.” *Logic Magazine*, August 31, 2020. https://logicmag.io/care/informatics-of-the-oppressed/.

O’Mara, Margaret. *Cities of Knowledge: Cold War Science and the Search for the Next Silicon Valley*. Princeton, NJ: Princeton University Press, 2004.

Parisi, David. *Archaeologies of Touch: Interfacing with Haptics from Electricity to Computing*. 1st edition. Minneapolis: Univ Of Minnesota Press, 2018.

Paterson, Mark. *Seeing with the Hands: Blindness, Vision and Touch After Descartes*. 1 edition. Edinburgh: Edinburgh University Press, 2016.

———. *The Senses of Touch: Haptics, Affects, and Technologies*. Oxford ; New York: Berg, 2007.

Paterson, Mark, and Martin Dodge, eds. *Touching Space, Placing Touch*. Farnham, Surrey, England ; Burlington, VT: Ashgate Pub, 2012.

Petrick, Elizabeth R. *Making Computers Accessible: Disability Rights and Digital Technology*. Baltimore, MD: Johns Hopkins University Press, 2015.

Pickering, Andrew. *The Mangle of Practice: Time, Agency, and Science*. Chicago, IL: University of Chicago Press, 1995.

Pink, Sarah. *Doing Sensory Ethnography*. London, UK: SAGE Publications Ltd, 2015.

Plotnick, Rachel. "At the Interface: The Case of the Electric Push Button, 1880–1923." *Technology and Culture* 53, no. 4 (November 20, 2012): 815–45. https://doi.org/10.1353/tech.2012.0138.

———. *Power Button: A History of Pleasure, Panic, and the Politics of Pushing*, 2018.

Riskin, Jessica. *Science in the Age of Sensibility: The Sentimental Empiricists of the French Enlightenment*. Chicago: University of Chicago Press, 2002.

Rosner, Daniela K. *Critical Fabulations: Reworking the Methods and Margins of Design*. Cambridge, MA: The MIT Press, 2018.

Scherffig, Lasse. "There Is No Interface (without a User): A Cybernetic Perspective on Interaction." Edited by Florian Hadler, Alice Soiné, and Daniel Irrgang. *Interface Critique Journal* 1 (2018): 58–80.

Semel, Beth M. "Listening Like a Computer: Attentional Tensions and Mechanized Care in Psychiatric Digital Phenotyping." *Science, Technology, & Human Values* 47, no. 2 (March 2022): 266–90. https://doi.org/10.1177/01622439211026371.

Shackel, B. "Man-Computer Interaction-The Contribution of the Human Sciences-," n.d., 15.

Shneiderman, Ben. "The Future of Interactive Systems and the Emergence of Direct Manipulation." *Behaviour & Information Technology* 1, no. 3 (July 1, 1982): 237–56. https://doi.org/10.1080/01449298208914450.

Small, Mario Luis, and Jessica McCrory Calarco. *Qualitative Literacy: A Guide to Evaluating Ethnographic and Interview Research*. Oakland, California: University of California Press, 2022.

Snow, C. P. *The Two Cultures*. Cambridge University Press, 2012.

Stigler, George J. "Economics: The Imperial Science?" *The Scandinavian Journal of Economics* 86, no. 3 (1984): 301–13. https://doi.org/10.2307/3439864.

Stokes, Dustin, Mohan Matthen, and Stephen Biggs, eds. *Perception and Its Modalities*. 1 edition. London ; New York, NY: Oxford University Press, 2014.

Stoller, Paul. *Sensuous Scholarship*. Philadelphia: University of Pennsylvania Press, 1997.

Suchman, Lucy A. *Plans and Situated Actions: The Problem of Human-Machine Communication*. Palo Alto, CA: Palo Alto Research Center, Xerox Corporation, 1985.

Sutherland, Ivan Edward. "Sketchpad: A Man-Machine Graphical Communication System." Ph.D. Thesis, Massachusetts Institute of Technology, 1963.

Tripodi, Francesca. "Searching for Alternative Facts." Media Manipulation. Data & Society Research Institute, n.d.

Velasco, Carlos, and Marianna Obrist. *Multisensory Experiences: Where the Senses Meet Technology*. Oxford: Oxford University Press, 2020. https://doi.org/10.1093/oso/9780198849629.001.0001.

Williamson, Bess. "Getting a Grip: Disability in American Industrial Design of the Late Twentieth Century." *Winterthur Portfolio* 46, no. 4 (December 2012): 213–36. https://doi.org/10.1086/669668.

Wright, Sarah H. "Professor Robert W. Mann, Leader in Prosthetics, Dies at 81." MIT News, June 19, 2006. https://news.mit.edu/2006/obit-mann.

Chapter 1

Let’s Go, Baby Forklift!: Fandom Governance and the Political Power of Cuteness in China

Abstract

This article describes how the Chinese state borrows from the culture of celebrity fandom to implement a novel strategy of governing that we term “fandom governance.” We illustrate how state-run social media employed fandom governance early in the COVID-19 pandemic when the country was convulsed with anxiety. As the state faced a crisis, state social media responded with a propagandistic display of state efficacy, broadcasting a round-the-clock livestream of a massive emergency hospital construction project. Chinese internet users playfully embellished imagery from the livestream. They unexpectedly transformed the construction vehicles into cute personified memes, with Baby Forklift and Baby Mud Barfer (a cement mixer) among the most popular. In turn, state social media strategically channeled this playful engagement in politically productive directions by resignifying the personified vehicles as celebrity idols. Combining social media studies with cultural and linguistic anthropology, we offer a processual account of the semiotic mediations involved in turning vehicles into memes, memes into idols, and citizens into fans. We show how, by embedding cute memes within modules of fandom management such as celebrity ranking lists, state social media rendered them artificially vulnerable to a fall in status. Fans, in turn, rallied around to “protect” these cute idols with small but significant acts of digital devotion and care, organizing themselves into fan circles and exhorting each other to vote. In elevating the memes to the status of celebrity idols, state social media thereby created a disposable pantheon of virtual avatars for the state, and consolidated state power around citizens’ voluntary response to vulnerability. We analyze fandom governance as a new development in the Chinese state’s long history of governing citizens through the management of emotion.

Publication: This chapter appears as Jamie Wong*, Crystal Lee*, Vesper Long*, Di Wu*, and Graham Jones*. “Let’s Go, Baby Forklift: The Power of Cuteness and the Cuteness of Power on Chinese Social Media,” *Social Media + Society* (April 2021).

Acknowledgements: The authors are grateful to Zizi Papacharissi and this journal’s anonymous reviewers for their transformative input, and to Yilia Qu, Bambi Schieffelin, Val Wang, Rob Weller, Mei Ling Young, and the members of the Language and Technology Lab for invalu- able feedback. They also thank Rachel Heiman and her students at the New School for their provocative engagement with a draft of this article.

When China began its massive mobilization against the novel coronavirus in January 2020, no one could have imagined that a little yellow forklift would come to embody the nation's hope. Nevertheless, by the end of the month, "Let's go, Baby Forklift!" had become a collective rallying cry, as millions watched an army of construction workers erect emergency hospitals in Wuhan, the pandemic's epicenter. In this article, we show that the little forklift could carry the burden of a nation's fears because it was not just any construction vehicle. The collective efforts of an adoring public and the guiding hand of Chinese state social media had transformed it into a much more powerful kind of *symbolic* vehicle: a celebrity idol. As an idol, it could command devotion not through strength alone, but also through vulnerability (Pritzker 2020)—a quality that made it cute. How did a burly machine become cute and why did cuteness make it such a potent national symbol during a moment of crisis? In answering these questions, we show how construction vehicle memes came to figure into the Chinese state's ongoing efforts to harness the iconography and interactional forms of Chinese internet fandom for political purposes.

For decades, scholars have been tracking the evolution of the Chinese state's strategies of managing public information and influencing opinion alongside the rapidly changing digital affordances of the social mediascape (Brady 2002). During the 2003 SARS outbreak, it deployed what Lagerkvist (2008) called "ideotainment," sending out messages and images of heroic figures to individual mobile phones across China to assure the public that the state had the crisis under control. In intervening years, the state has developed a vast apparatus of paid social media monitors—the "Fifty-Cent Army"—to shape public opinion, while also fomenting "voluntary" networks of nationalistic content-producers (Han 2015a). Han demonstrates that resulting social media influence campaigns are not, however, a "centrally coordinated process in which the state acts as a monolithic rational actor." Instead," he continues, "multiple party-state agencies at different levels and in different sectors as well as individual officials are involved, each with different incentives and priorities" (Han 2015b, p. 109). They can, in turn, learn from, adopt, and adapt each other's innovative strategies.

COVID-19 put the function of this propaganda and censorship system on stark display, while also revealing some of its limitations. During the first months of the pandemic, regional governments used sophisticated microwork platforms to coordinate what may have amounted to "hundreds of thousands" of commenters and social media monitors tasked with posting original content, sharing links, and flagging negative comments (Zhong et al. 2020). In this article, we offer an archival ethnography of a digital "media event" (Pink et al. 2015) that unfolded within this social media ecology, at the intersection of state influence and public participation. At a time when much of China was under quarantine and the country was convulsed by anxiety about the novel coronavirus, we show how Chinese internet users playfully embellished images of the hospital construction broadcast online by state-run social media, producing memes that were, in turn, assimilated back into state-run social media feeds in an emergent cycle of cultural co-production.

It is important to stress at the outset that millions of social media users were directly and indirectly involved in generating the memes we analyze. We do not and cannot claim to directly understand their intentions, nor to accurately gauge the involvement of paid influencers. Social media platforms, Christopher Kelty (2020) writes, make it possible "to participate in something and be given signs and signals of the outcomes of that participation" (p. 257), offering "a set of scripts and games intended to spread participation as widely as possible" (p. 261). Thus, what we study here is how memes function as publicly visible *signs of participation* in a game-like sequence where participants signal involvement by responding to each other's "moves" and making their responses available for others to scrutinize and incorporate into their own subsequent moves. In developing our analysis, we drew upon one co-author's firsthand knowledge of these memes, which began popping up in her social media feeds as the events described here unfolded. We constructed our corpus by taking pivotal hashtags—signs of participation that contributed to shaping the "eventfulness" of this episode—as focal ethnographic objects (Bonilla and Rosa 2015), guiding the collection of publicly available content. Finally, in interpreting the complex sociocultural knowledge embedded in linguistic and visual signs, we elicited additional feedback from Chinese ethnographic respondents.

This approach enables us to identify how Chinese state social media accounts actively mobilized the game-like repertoire of celebrity fandom to produce public consensus through seduction rather than coercion. As W. Zhang (2016) argues, the Chinese entertainment industry relies on social media infrastructure that seductively stages the passionate participation of "fandom publics" as an integral part of celebrity idols' allure. That infrastructure includes ludic modules such as ranking lists that organize fans into games of voting competitively for their favorite idols and provide real-time participatory feedback (Yin 2020; Q. Zhang & Negus, 2020). We show how Chinese state social media accounts deployed such game-like features as scripts to enlist fandom publics around the construction vehicles and make fans' collective performances visible. We demonstrate how ludic modules for managing fans became tools for governing citizens-as-fans, as the state used them to generate viral "positive energy" across internet platforms during an overwhelmingly dark time. Our analysis focuses solely on social media content and strategies, not the state's policies to contain COVID-19.

We combine insights from social media studies with cultural and linguistic anthropology to offer a processual account of the semiotic mediations involved in turning vehicles into memes, memes into idols, and citizens into fans. In particular, we focus on cuteness as the crucial semiotic resource in the co-construction of seductive memetic idols and an adoring public of devoted fans. Few people would describe any state as "cute," and the vastly powerful Chinese state may seem particularly un-cute. Nevertheless, in the events we analyze, cuteness emerges as a dominant motif in Chinese citizens' engagements with the state, and in the state's engagement with citizens. We argue that cuteness, in conveying qualities of dependency and diminutiveness, can constitute a counterintuitive form of power: the power to inspire affection and to engender play (Abidin 2016; Hiramoto & Wee 2019; Silvio 2019). In tracing how Chinese internet publics and state-run social media accounts interactively mobilized signs of cuteness during a national crisis, we offer an account of both the power of cuteness and the cuteness of power in contemporary Chinese society.

In our case study, Chinese internet users playfully manipulate imagery from a state-run livestream during a pandemic to produce memes in ways that might, at first blush, seem irreverent or subversive. However, we argue instead that playfulness is part of an emerging pattern of *ludic governmentality* in state-run social media influence campaigns. By "governmentality," Foucault (1991) points to forms of social control exceeding the bounds of "government" in the strict sense of political institutions. Anthropologists have shown that the regimentation of play can be a profound manifestation of governmentality (Walker 2013). With the notion of ludic governmentality, we point to a relationship, on necessarily unequal terms, between state and subjects as playmates. It is commonly assumed that authoritarian regimes derive power through monologic control over the stable meanings of political discourse and national symbols. While it follows that playful memetic content can be a powerful form of resistance (Johnson 2015), we show how the calculated introduction of framing scripts—here as a form of play-based governance—can generate a controlled polyphony that reinforces state authority.

In the episode we analyze, by capitalizing on cuteness to engage internet users in online play, official accounts strategically shifted discourse about an unfolding crisis. As we demonstrate, when internet users acknowledged state social media's effort to engage with them in play, they interpreted it in contradictory terms. Where some saw benevolent care and paternal affection, others saw crass pandering and inappropriate manipulation. This divergence reflects the way play itself ambiguates structure and agency (Sutton-Smith 1997), particularly when participants freely elect to follow arbitrary rules set by someone else. Focusing on critiques of the construction vehicle memes in the final section of this article, we explore the kinds of risks—slippages in meaning, divergent interpretations—the state runs when deploying the ludic repertoire of social media fandom as a mode of governance.

1 From Spectators to Supervisors

Following the initial appearance of COVID-19 in Wuhan in December 2019, rumors about the virus—and criticisms about the official response—circulated widely on Weibo, WeChat, and other social media platforms (Weinland 2020). Wuhan residents posted videos showing the city's hospitals overwhelmed and ill-equipped to deal with the pandemic, provoking outrage about mismanagement within the Red Cross, a humanitarian organ of the Chinese state (Yuan 2020). With reports of the first casualties and cases appearing in other provinces, the central government intervened to assert control over the flow of information (Wang 2020). On 23 January, it placed nearly 50 million people in Hubei province on lockdown, creating what international observers heralded as “the largest quarantine in human history” (Ang 2020). That same day, the government began constructing two emergency field hospitals in Wuhan: the 1000-bed Huoshenshan (火神山, Fire God Mountain) and 1600-bed Leishenshan (雷神山, Thunder God Mountain). In a stunning display of efficiency, they would be operational in just over a week (Qin 2020).

On 27 January, the internet video platform Yangshipin (央视频), which is under the control of the Publicity Department of the Chinese Communist Party, announced that it would offer a round-the-clock livestream of the construction project, inviting internet users to “witness” (见证) the colossal undertaking for themselves. As construction workers toiled day and night, rapt citizens logged on by the tens of millions to watch aerial views of colorful construction vehicles bustling back and forth. Viewers posted supportive comments that scrolled beneath the video in real time. As signs of engagement, these comments helped turn mass viewing itself into a spectacle of participation (cf. Cao 2021). In an apparently sui generis display of playful involvement, some of the commenters began referring to themselves as “supervisors” (监工). “Henan supervisor has arrived,” one announced. “Local Hubei supervisors are here to take over. Colleagues, go eat,” another chimed in. Within 24 hours, commenters were describing themselves as “internet supervisors” (网络监工) and “cloud supervisors” (云监工). After CCTV’s Weibo account began using the hashtag #InternetSupervisors (#网络监工#), “remote supervision” became a widespread trope in posts about the livestream from internet users and governmental agencies alike, amplifying publicity. By 30 January, the livestream had already received 190 million views and 77,000 comments.

Livestreaming the construction project may have been a relatively straightforward way for state social media to counter a crisis of public confidence with a meta-message of transparency. However, by shifting the idiom of viewership from “witnessing” to “supervising,” commenters introduced a more active model of participation that ultimately opened the door to fandom through play. As we will show, this humorous online banter was the first in a series of “play signals” (Bateson 2000) that catalyzed increasingly ludic forms of interaction and mobilized a detached public into devoted fans, whose visible participation was central to the spectacle’s appeal. In the course of several days, these “remote supervisors” produced elaborate memes based on the personification of construction vehicles that they anointed with cute and humorous names. Ultimately, state social media would deploy games and scripts that created the conditions for remote supervisors to organize themselves into fan groups supporting these idolized heroes, who came to serve as embodiments of the nation.

The livestream of the hospital construction project brought into relief Chinese citizens’ wholesale dependence on the state. As Ang (2020) notes, after years of intensified centralization had gutted civic sectors, the hospital construction project dramatized the kind of mass mobilization possible *only* under the auspices of the national government. Yet this staggering spectacle of state efficacy came to derive much of its power from dialogism:

internet users' active engagement as not just patriotic citizens—but also playful supervisors and devoted fans—would make their participation a crucial component of its propagandistic appeal. This is symptomatic of the Chinese government's increasing investment in what Repnikova and Fang (2018) call "a more participatory form . . . of official persuasion online" (p. 770). By tracing the genesis of playful memes against the backdrop of a crisis, we show how the cultural repertoire of participatory fandom intersected with propagandistic state social media to create user-generated, nationalistic content: "participatory propaganda" (Asmolov 2019). Just as fans' devotion cannot be mandated but must be given voluntarily, the most effective forms of propaganda are not manufactured by the state, but generated in a participatory fashion by citizens themselves.

2 Conceptual Framework: Fandom Governance

Scholars of the Chinese internet have charted the convergence of social media fandom and cyber-nationalism in what Liu (2019, p. 141) calls "fandom nationalism." Recent scholarship has focused on its manifestations in outward-facing foreign "campaigns" or "expeditions," in which Chinese internet users employ the tactics of Chinese fandom communities to coordinate social media attacks against political adversaries. For instance, in 2017, internet nationalists using the feminized moniker "Little Pink" inundated Tsai Ing-wen's Facebook page with pro-China slogans in an event now called the Diba Expedition (帝吧出征), catapulting fandom nationalism to widespread attention (Fang & Repnikova 2018).

There are several ways of understanding how the Chinese state foments and benefits from fandom nationalism, broadly conceived. On the simplest level, authorities can enlist internet users to generate online distractions that drown out disapproved topics (Moskowitz 2019). Guobing Yang points to a more subtle way in which the state cultivates the tenor of online discourse. Besides "repressive" forms of censorship that remove offensive content from circulation, Yang (2019) notes that it also engages in "productive" approaches by encouraging user-generated expressions of nationalism and amplifying them whenever possible (p. 9). Similarly, in our case study, state social media's amplification of the construction vehicle memes appears to drown out dissent with coordinated, politically productive content (Roberts 2018).

Building upon this scholarship, we draw attention to what we call *fandom governance*. We study how, during a moment of national crisis, state-run social media adopted a strategy of governing citizens-as-fans by borrowing registers (Agha, 1999), scripts, and games (Kelty 2020) from Chinese internet culture around reality television talent competitions. In doing so, we investigate state-run social media's role in setting the terms of participation on social media platforms (Gorwa 2019) to promote worshiping the nation according to the conventions of worshiping celebrity idols (Liu 2019).

Fandom governance fits into China's "longstanding history" of concern with "the effective harnessing of emotions for the sake of governance" (Steinmüller 2015, p. 3): from Imperial China to the Communist Era, the Chinese state has long used filial piety toward patriarchal figures as a model for subjects' devotion to rulers. More recently, as postsocialist Chinese families have come to value loving intimacy over formal hierarchy, the state also aspires to make itself loveable: as "love becomes a common language available to the private sphere, . . . it is also elevated to the level of propaganda in order to encourage and strengthen people's devotion to the regime" (Guo 2020, p. 48). In our case study, we see modalities of both filial and intimate devotion, with the critical addition of cuteness as a resource in the state's productive management of emotion.

We extend prior genealogical mappings of Chinese cultural models of devotion and affect management by drawing attention to the notion of *meng* (萌), an adjective meaning "cute" that originated in the Chinese internet slang before spreading into common usage (de Seta 2014). It is a loanword derived from the Japanese *moe*, which refers to affection, possessiveness, and protectiveness of the sort that fans feel toward animated characters or pop culture icons. As a vernacular form of East Asian cuteness cultures detailed by Abidin (2016), *meng* remains closely associated with the subculture of Chinese internet fandom from which it emerged, and has spawned compound words describing the practices of fans and idols. The verb *menghua* (萌化), literally "to cutify," means making someone or something cute so that it can be adored or even idolized. The verb *maimeng* (卖萌), literally "to sell cuteness," means passing oneself off as cute to inspire adoration, leading us to opt for a more liberal translation: "playing cute." Within the playful arena of internet fandom, fans and idols engage in co-constitutive moves: fans cutify idols and idols play cute. A fan's vulnerability to *meng* is linked to the inherent paradox of neoliberal agency: the assertion of individualistic selfhood by succumbing to the entertainment industry's consumerist desires (see Ngai 2012; Rofel 2007).

We demonstrate that the assimilation of the field hospital construction project to the culture of internet fandom hinged on the affective potential of *meng*: remote supervisors cutified (*menghua*) construction vehicles; state social media, in response, played cute (*maimeng*), connecting fans' affection for the vehicles with the state's desire for citizens' affection. By tracing manifestations of *meng* in the construction vehicle memes, we provide an ethnographic account of this powerful new principle of relational intimacy in contemporary Chinese culture. In the episode we analyze, *meng* functioned to create "scalar intimacies" (Pritzker & Perrino 2020) by associating different types of relationships—parents and children, fans and idols, and citizens and the nation—through verbal conventions of interpersonal affection. As remote supervisors evolved into fans, they were able to enact these connections across multiple scales thanks to a broader communicational "infrastructure of intimacy" (Wilson 2015) encompassing social media platforms such as Weibo and the state-run social media accounts that leapt in to capitalize on the opportune emergence of cute memes with patriotic shadings. Although endemic to the arena of entertainment, the playful sociability associated with *meng* proved wildly extendible: cutification was a means of constituting the vehicles as celebrity idols, while playing cute was a means for the state to court adoration.

This episode figures into the Chinese state's ongoing efforts to successfully harness the power of cuteness and its potential to politically engage fandom publics through social media. As early as 2013, a Weibo account personifying the Chinese space agency's lunar rover as a cute bunny (萌兔) attracted almost a million followers. The Chinese Communist Party's (CCP) official media outlet called upon other government agencies to follow suit by "playing cute" (*maimeng*): personifying themselves, titillating emotions, and capitalizing on signs of public engagement (Wu 2016)—precisely the playbook that unfolds in our case study. However, a number of later attempts to manufacture idols proved futile. A 2019 propaganda anime series commissioned by the CCP presenting a svelte young Marx as an idol met with mixed reviews (Douban 2019). A year later, the CCP Youth League's personifications of itself as two "virtual idols"—the girl Tender Country (江山娇) and the boy Blazing Flag (红旗漫)—proved a flop and were taken down within hours due to social media backlash (Phoenix Media 2020). The contrast between these successes and failures points to a central feature of fandom governance. While the CCP has long been attuned to the potential of marshaling fandom for political mobilization, being able to do it successfully seems to require allocating fans themselves some measure of what Kelty (2020, p. 14) terms "contributory autonomy": opportunities for participants' involvement in creative co-production.

Linguistic anthropologist Teri Silvio (2010) influentially theorizes the participatory investment fans make in popular East Asian characters as a form of "animation." For Silvio (2010), animation is "the projection of qualities perceived as human—life, power, agency, will, personality, and so on—outside of the self, and into the sensory environment, through acts of creation, perception, and interaction" (p. 427). As a way of thinking about how personified characters at the heart of anime, manga, and gaming industries come to life, this framework is generative on multiple levels. First, it highlights that these characters are the "creations of collectives" (p. 428), involving a complex division of labor between professionals with a variety of interlocking skills. Perhaps more importantly, "fans, and many scholars, often . . . sense that these characters have lives of their own . . . arising from their re-creation in numerous media and styles by hundreds, thousands of fans" (p. 428). Given the power of this framework for mobilizing emotional investment, Silvio argues, it has spread from the realm of entertainment; now corporations, nations, and even religious movements seek to brand themselves with cute mascots capable of inspiring devoted followings (Silvio, 2019). We build upon Silvio's insights, showing how the co-created animation of the construction vehicle idols draws upon the cultural logic of fandom. It hinges on an incompleteness that state social media might seek to canalize but can never fully control.

In the following sections, we examine how an incipient fandom takes shape around the field hospital construction vehicles through the process of animation. We focus on the personifying practices of naming the construction vehicles, visually depicting their identities, and mobilizing them as interactants. We note that because this process of animation is collective and distributed across myriad animators with distinctive perspectives, the personified vehicles have multiple names and protean identities. We show that, as animated vehicles took on life as the objects of remote supervisors' adoration, Chinese state social media accounts channelled the outpouring of emotion into the participant framework and digital infrastructure of Chinese fandom culture. In capitalizing on the popularity of the construction vehicles and consolidating interpretations around themes of national unity through fandom governance, the state became an active participant in the animation process.

3 Animation Through Naming

On the second day of the livestreaming, the participant structure of remote supervision coalesced alongside the animation of construction vehicles as personified figures. Naming was central to the process of animation, individuating the vehicles and contributing to their portrayal as characters. In the comment section under the livestreams and in posts on Weibo, remote supervisors began referring to the vehicles using affectionate nicknames in the way one might address a child. Nicknaming may have begun spontaneously, but it eventually became an activity in its own right, and a form of verbal play. "Chinese nicknames have as their primary function the creation or validation of a sense of casual, good-natured fun," explains Moore (1993, p. 75). The kinds of affectionate diminutives remote supervisors gave to the construction vehicles were analogous in structure and function to childhood "family nicknames" (小名, literally "little name") that "define membership in the kinship group" by expressing "affectionate intimacy" and "a very strong sense of family loyalty" (p. 80). Thus, diminutive nicknames served to simultaneously "key" (Goffman 1974, p. 43) not just play per se, but the particular kind of playfulness associated with domestic intimacy. This intimacy is reflected in the remote supervisors' comments. "I'm watching Huoshenshan Hospital grow up," one wrote. Here, the character for "watch"—*kan* (看)—conjoins different modalities of care. It refers at once to witnessing as a concerned spectator, monitoring as a remote supervisor, and looking after one's children as a parent. From the use of diminutives and explicit reference to child care and rearing (看着长大), fictive kinship emerged as a vital element in the transformation of this event into an instance of fandom nationalism highlighting shared membership in a "national family" (国家).

Some remote supervisors formed diminutives by conjoining the prefix "Little" (小) with the vehicle type, as in Little Forklift (小叉车), or with a distinguishing color attribute, such as Little Yellow (小黄, a mechanical digger) or Little Green (小绿, a dump truck). The reduplication of the diminutive could add an extra-cute effect: Wee Little Yellow (小小黄). Closely related terms of endearment featuring the prefix "Big" (大) also appeared in the comments, for instance, Big Yellow (大黄, a pump truck). In Chinese, such augmentative nicknames function according to a logic similar to diminutives: the prefix "little" conveys the same affectionate intimacy toward status subordinates as the prefix "big" toward status superordinates. Generally speaking, all these terms of endearment were mechanisms of *menghua* that enabled remote supervisors to animate the vehicles as personae, imbue them with qualities of cuteness, and express affection toward them.

In a pattern that merits particular attention, remote supervisors formed additional nicknames by appending "baby" (酱, *jiang*)—an ACGN (Anime, Comics, Video games, and pulp Novels) internet slang suffix deriving from the Japanese diminutive "-chan" (ちゃん)—to some distinguishing attribute of a particular vehicle: Baby Forklift (叉酱), Baby Tractor Shovel (铲酱), Baby Red (红酱), Baby Green (绿酱), and so on. For added cuteness, some even combined both a diminutive prefix and suffix, for example: Baby Little Red (小红酱) or Baby Little Green (小绿酱). Pointing indexically to both the Japanese ideal of *kawaii* cuteness (Hiramoto & Wee 2019) and the Chinese internet fandoms that draw upon that cultural repertoire, suffixation with *jiang* conveys endearing infantilization, which is why we gloss it as "baby." Consequently, the diminutive appellations formed with this suffix suggest a feminized form of affection—the type of care and concern one directs to an infant—while also evoking the feminized subculture of fan circles (饭圈), where it is commonly used to create discursive intimacy with idolized characters and celebrities. Underlining this association with internet fandom, a popular cement mixing truck earned the appellation Baby Mud Barfer (呕泥酱, *ounijiang*), which sounds like an affectionate Japanese term for "big brother" (お兄さん, *oniichan*), widely used among young Chinese participants in ACGN online subcultures.

Because these names emerged spontaneously from the engagement of an incipient fandom, many vehicles received multiple names, sometimes with strikingly different connotations. Despite the prevalence of diminutive nicknames drawing on the feminized idiom of cuteness, remote supervisors also coined nicknames evocative of a masculine, warrior ethos. Several names were homophones for the titles of historical figures, especially emperors. Besides Baby Mud Barfer, remote supervisors called the cement mixer Emperor Who Sends Dust (送灰宗), a homophone for of Emperor Huizong (宋徽宗) from the Song Dynasty. They came to call a blue-colored mechanical digger Emperor Who Always Works (勤史皇), a homophone for the first emperor of a unified Imperial China, Qin Shi Huang (秦始皇). A vehicle involved in electric welding earned the appellation Welding Emperor (焊舞帝), a homophone for the powerful Emperor Wu (汉武帝) of the Han Dynasty. As we show below, these historic nicknames—in contrast to the cutesy nicknames—enabled different interactions between remote supervisors and the idols they worship.

These ingenious homophones draw on the playful use of sound-alike terms in Chinese internet culture, and particularly humorous memes. While homophones are otherwise often associated with flouting state censorship (Miltner 2018; Zidani 2018), here they index powerful male protectors and progenitors from Chinese history to create provocative analogies between past and present. As Moskowitz (2019) writes of the place of historical allusions in Chinese internet culture, “slamming history and the present moment together in this fashion is a profound revisualization of both historical events and the contemporary era,” creating both playful juxtapositions and serious evocations of an inspirational tradition (p. 20). Thus, one remote supervisor posted a comment describing historical homophones as “the upgraded (升级) versions” of the cute terms of endearment that initially appeared in comment threads. These “more high-end and magnificent” (高大上) names, the comment continued, “give us a feeling that our ancestors from every dynasty are protecting us.” Significantly, this comment very clearly distills the paternalistic model of ancestral worship of patriarchal protectors and implicitly extends it to the contemporary Chinese state.

It is striking that two paradigms of nicknames emerged drawing diametrically on the domain of kinship, one inserting the vehicles into the role of small children who need parental protection, the other positioning them in the role of ancestors protecting living descendants. By the time remote supervisors began circulating compendia of vehicle names 2 days after the start of streaming, some vehicles had multiple nicknames based on different naming conventions. For instance, the digger named Emperor Who Always Works (勤史皇) could also be called Little Blue (小蓝) or Lan Wangji (蓝忘机), a homophone for the hero in China’s then most popular television series, *The Untamed*. The multiplicity of names points to a key feature of playfulness in internet culture: creation is always partial, incomplete, open-ended, and available for remix (Miltner 2018; Moskowitz 2019; Phillips & Milner 2017). The names thrive on polysemy and associations across cultural strata ranging from Japanese ACGN, ancient imperial history, to contemporary television shows.

These cute nicknames also drew intertextually on other recent Chinese memes, notably those associated with fandom nationalism. In 2019, participants in a social media campaign against Hong Kong protestors proclaimed their intent to defend Older Brother A-Zhong (阿中哥哥), a personification of the Chinese nation in the form of a celebrity idol (Initium Media 2019). Moreover, the construction vehicle nicknames came to figure centrally into a pattern of giving cute nicknames to other inanimate entities associated with the COVID-19 pandemic. Thus, in online discourse, the virus itself became Little Corona (阿冠) and Cute Sister (萌妹), a charming but dangerous seductress. Wuhan also garnered two novel nicknames: Little Fool (小笨蛋), for its ignominious role in the pandemic, and Hot-Dry Noodles (热干面), a famous regional dish.

4 Animation Through Visualization

Weibo users added to the identity of each vehicle by producing image macro memes (Miltner 2018; Mina 2019) with vehicles’ nicknames superimposed on screen grabs from the video feed. Individual users and state social media accounts widely circulated these memes separately and in curated galleries. For instance, Figure 1 shows a gallery featuring several images of the cement mixer Baby Mud Barfer, with memes depicting its various nicknames: Emperor who Sends Dust, Big White Rabbit (大白兔, which is also the name of a favorite children’s candy), and White Chubby-Chub (白滚滚). Some of these image macros added an additional layer of animation by personifying vehicles' behavior. For instance, the meme in Figure 2 shows two of these cement mixers parked together with their cabins angled toward each other. The caption reads "Baby Mud Barfers whispering together" (窃窃私语的呕泥酱). Animation here is not just a matter of imputing personalities to the construction vehicles, but also playfully positing storylines that suggest social interactions and ongoing relationships between them.

Figure 1. Weibo post from Yangshipin sharing a curated gallery of vehicle nicknames and memes.

Figure 2. Meme anthropomorphizing "Baby Mud Barfers whispering together."

Figure 3. Baby Forklift "burden of life" meme.

Personifying the vehicles was also a means of identifying with them: in the meme depicted in Figure 3, an especially popular vehicle, Baby Forklift, carries two gargantuan bent pipes. The superimposed text reads "the burdens of life" (生活的重担), a catchphrase used in Chinese social media to joke about everyday hardships. Intertextual references like these establish associations with Chinese internet culture, and generate humor by creatively redeploying familiar laugh-lines. Memetic humor presumes virtual relationships based on shared sentiment: memes "promote a specific form of internet literacy that is embedded in virtual social networks and geared towards the affective response of laughter—a social practice of laughing with, as well as laughing at" (Rea 2013, p. 171). Thus, the proliferation of playful memes and humorous content around the construction vehicles reflects a virtualized form of collective effervescence based on shared recognition of cultural and subcultural intertexts, and shared experiences of pleasure (Shifman 2013). This event clearly corroborates Coates's (2017) observation that the effervescent virality of memes in Chinese social media reflects local "understandings of sociality as spreadable, seductive, and affective" (p. 163).

In addition to generating and circulating image macros, other participants contributed to the visual personification of the vehicles through extraordinary pieces of original artwork. Figure 4 is an image personifying Baby Mud Barfer in the style of *doujin* (同人, an ACGN genre often associated with fanart). Although the vehicle's nickname is splashed across the background, the cement mixer itself is only partially visible. The artist has personified the vehicle in the figure of its driver, portrayed according to the aesthetic conventions of "soft masculinity" (Jung 2010). Closely associated with East Asian pop idols and celebrity boy bands, this kind of masculinity emphasizes male beauty and emotional sensitivity. Soft masculinity is manifest in several ways here, beginning with the driver's fashionable and futuristic attire, which is immaculately clean. With large, lachrymose eyes, he is a study in melancholy. Moreover, note the juxtaposition between his slight, androgynous frame and the strength of his workmanly hands. Overall, the image evokes a conceptual contrast between individual human frailty and the weight of historical responsibility, doubly personified through the paired figures of Baby Mud Barfer and its driver. The depiction of the driver as an object of romantic contemplation according to visual conventions of celebrity idol worship gives the image an almost votive quality.

Figure 4. Fanart anthropomorphizing Baby Mud Barfer as a "soft" male idol.

The images we have described so far all involve aspects of cutification. With no overt reference to the context of the pandemic, they soften the seriousness of the vehicles and turn them into objects of play and affection. Other user-generated imagery, however, embellished the martial associations reflected in the historical nicknames, foregrounding traditionally "hard" masculinity. The artwork shown in Figure 5 offers one vivid example, drawing heavily on the visual vocabulary of anime and manga. It shows some of the most popular construction vehicles, Baby Forklift in the lead, wearing surgical masks and arrayed heroically in front of one of the construction site's white-colored temporary buildings, with each of their nicknames written overhead. The Gods of Fire and Thunder, for whom the field hospitals are named, tower above them, menacing immortals with rippling muscles. The words at the middle of the picture (捍武大帝) are homophones for the title of one of ancient China's most powerful Emperors, "The Great Conqueror King of Han" (汉武大帝). The substituted first character (捍) means “protect,” suggesting a new signification for the second character (武) “Wu(han),” and yielding a new meaning: “Great King Protector of Wuhan.” Images like this animate the cute vehicles through the patriarchal idiom of ancestral worship. In doing so, they position the accelerated construction of hospitals as continuing the dynastic tradition of safeguarding the Chinese state and its people, lending an air of historicity and monumentality to the vehicles.

Figure 5. Fan-drawn poster depicting construction vehicles as “hard” masculine heroes in an epic battle.

5 From Memes to Idols: Animation Through Fandom Governance

Naming and visual depiction were crucial parts of animating the vehicle memes as characters, but elevating them to the status of celebrity idols involved “scripting” (Kelty 2020, p. 261) participants as fans. By introducing ludic modules that channeled remote supervisors’ adoration into systems of voting and ranking, state social media made their contributory involvement visible as idol worship. In the context of Chinese fandom, symbolic exchanges between fans and idols are public displays of devotedness that culturally magnify the idol's persona. Ultimately these "parasocial" interactions can take the form of "parakin" relationships that require fans to "protect" and "support" their idols like family members (Yan & Yang 2020). Interacting with idols also creates "scalar intimacies" (Pritzker & Perrino 2020) as fans interact with each other in fan circles. In this way, remote supervisors came to socialize both *with* and *through* the construction vehicles, using cuteness as a resource for constituting intimacies through playful forms of involvement.

Viewers of the construction livestreams made comments that clearly connected the forms of affectionate care that supervisors show their workforce and that fans show their celebrity idols. For instance, a Weibo user posted, "Baby Tractor Shovel goes to bulldoze, supervisors follow you forever. To cheer on Little Yellow shout #GoForItWuhan." Another Weibo user posted:

> *Although I have to stay at home during the spring festival, I still have a job to do. As a remote supervisor, I will support and supervise workers and be grateful for their contribution. [. . .] Baby Red, Baby Green, Baby Tractor Shovel, Baby Little Red, Baby Little Green, let's do it.*

Comments such as these reveal the interactive premise of animation: the cuteness of the idols engenders protective "support" that brings them to life as participants in symbolic exchanges and parakin relationships with fans.

State social media took an active role in amplifying the conceptualization of viewership as "remote supervision" by encouraging its rapid evolution into a form of fandom. It is not surprising, in the context of a national crisis, that state social media would throw its weight behind online activity generating enthusiastic publicity for a reassuring spectacle of governmental efficacy (see Feng 2020). In this sense, the spread of memes related to the construction vehicles came to resemble what Palmer (2007) terms a "political fever," when informal signals "sent from above correspond with, open the space for, and amplify popular desire, which appropriates these spaces in unexpected ways, simultaneously complying with, [...] disrupting, and mirroring the projects of state hegemony" (p. 162). State-run social media accounts' advertisement of memes and adoption of popular language used in fan culture was an informal signal, notifying internet users that the state approved and encouraged the playful narrative about the national project.

When the official Weibo account of CCTV News made the post shown in Figure 1 on the third day of live streaming, it included the following comment alongside the gallery of memes:

> *Too cute! #TenMillionCloudSupervisorsIntelligence 😜Remote supervisors who felt bored but couldn't fall asleep watched the hospital construction livestream and gave names to vehicles in the field: Little Red, Little Green, Little Yellow, Baby Forklift, Baby Tractor Shovel, Baby Mud Barfer…Do you know who they are?*

This extraordinary post shows the dawning attunement of state social media to the productive political potential of the vehicle memes. CCTV introduced a hashtag here (#TenMillionCloudSupervisorsIntelligence) that naturalized the activity of "cloud supervision," highlighted the massive scale of participation, and celebrated the crowdsourced creativity of animating the construction vehicles. The hashtag is framed on one end with the interjection "too cute!" (太可爱), and on the other with an emoji conveying a sense of fun and zaniness. With winking approbation, the state recognizes and partakes in remote supervisors' spirit of playfulness. This message is driven home by the challenge to internet users to learn the vehicles' cute nicknames.

A wide coterie of state-controlled social media accounts joined in reposting memes and circulating lists of vehicle names. Besides CCTV, our corpus includes Weibo posts by People's Daily, News China, China National Radio, The China Youth Daily, the Central Committee of the Communist Youth League, and others. The Assets Supervision and Administration Commission of the State Council exhorted Weibo users:

> *Darlings, wake up, let's continue remote supervision of the construction project. Little Yellow, Little Blue, Big White, Emperor Who Sends Dust, Emperor Who Always Works... You need to be familiar with these names, otherwise you will be derided by others in the comment section.*

State social media's active promotion of the nicknames was an important step toward consolidating politically productive play around vehicles-as-celebrities.

In the days that followed, state social media took an active hand in transforming these characters into celebrity idols and, ultimately, organizing this play into a form of fandom governance. On 29 January, People's Daily if not solely created, then at least almost single-handedly popularized, a new hashtag #ConstructionVehicleIdolsGroup (#挖掘机天团#), which quickly went viral on Weibo. Using the phrase "idols group" (天团)—a term commonly used to refer to boy bands and girl groups—framed the construction vehicles as celebrity entertainers and remote supervisors, by extension, as fans. The post received more than 279,000 likes, thousands of reposts and comments, and echoes from multiple other state social media accounts. As the hashtag became popular, the livestream soared to over 40 million simultaneous views, and the comment threads were filled with speculation about which vehicles would *chudao* (出道) next. *Chudao* is a term that denotes the debut of performance artists onto the public stage to launch their careers. By employing a term associated with the entertainment industry, internet users further elaborated linguistic associations between the vehicles memes and celebrity idol worship—an indication of their receptiveness to the role assigned to them in the script of fandom governance.

It should not come as a surprise that state social media displayed a high degree of proficiency in mobilizing the participatory frameworks of fandom. In contemporary China, the state is deeply involved in shaping synergies between entertainment, advertising, and politics (Sullivan & Kehoe 2019). Not only did state-run social media adopt a socioculturally distinctive linguistic register (Agha 1999) gesturing at Chinese fandom practices, it also introduced participatory modalities that borrowed directly from TV networks' and internet platforms' fandom management playbook. One participatory modality was particularly important in turning the construction vehicle memes into a full-fledged instance of celebrity idol worship: voting systems.

Mirroring the rankings of celebrity idols, Yangshipin, the CCTV affiliate livestreaming the hospital construction, soon incorporated an interactive ranking list for construction vehicle celebrities under its video feed, encouraging remote supervisors to cast votes supporting their favorites (Figure 6). For fans, casting votes to keep an idol at the top of a ranking list is a form of protection against elimination in reality shows and other platforms of visibility (Yin 2020; Q. Zhang & Negus 2020). Unlike idol candidates on reality shows, the vehicles' presence on the construction site certainly did not hinge on fans' ardor. Yet by casting these cute memes as idol contestants, state social media rendered them *artificially* vulnerable to descending in status or acclaim. In the process, it created a small and disposable pantheon of virtual avatars for the state, consolidating the power of fandom governance by mobilizing the public's engagement around images of vulnerability. Confirming this point, as construction vehicles entered into the ambit of fandom, this interactive ranking list and other posts made by state-controlled media increasingly employed their cute nicknames, not the patriarchal, historical appellations—suggesting a deliberate selection of *meng* as a paradigm for fandom nationalism.

Heralding the advent of Yangshipin's ranking list, National Business Daily raved on Weibo that

> *the Construction Vehicle Idols Group . . . has now officially debuted! The live video stream is so intoxicating (上头) that [internet users] cannot help but indulge themselves. Become a remote supervisor and clock in every day. When you* puff the list *for the idols group, don't forget to pay attention to real-time updates on the epidemic.*

This post not only refers to fandom in terms of compulsively intoxicating affective states, but also enacts a verbal register distinctive of fan culture: in their vernacular, fans refer to casting votes to help their favorite celebrities ascend the popularity rankings as "puffing the list" (打榜). This post's gentle nudge to pay attention to the pandemic points to the tension between escapist fandom and the more serious imperative of keeping the public informed.

Figure 6. Construction livestream on Yangshipin platform with recently introduced voting system for vehicles: Baby Forklift tops the ranking with 189,760 votes, followed by Little Blue and Baby Mud Barfer.

The level of convergence among state-run social media accounts around this moment of fandom nationalism suggests, if not a concerted effort, then at the very least a deeply shared attunement to fandom's political potential (cf. Han 2015b). Remote supervisors also reflected explicitly on the state's deep involvement in their online activities. One fandom opinion leader commented that by "opening a ranking list for vehicles," state social media showed that "they really know how to play the game, hahahaha." The notion of "knowing how to play the game" (太会玩了) here can connote both having a good time and, more insidiously, deftly manipulating rules to one's advantage. In the thread below this post, the most upvoted comment stated, "official certification hahahahaha." Another responded, "public officials (官方)

work hard so that we won't be bored." Other comments proclaimed that "Papa CCTV (央视爸爸) is comforting children" or that "officials are comforting us, the children who feel bored at home." In these comments, remote supervisors constitute state social media's deployment of fandom practices as another Batesonian play signal. Participants acknowledge the state for engaging subjects through ludic governmentality—distracting them from their fears, as a parent might do for a child. Yet the use of phrases like "Papa CCTV" highlights the imbalance of power implicit in this relationship, and the signs of laughter in several of these comments lend themselves to a variety of interpretations, ranging from genuine delight to dismissive cynicism. Taken together, these comments display recognition that, in amplifying animation through practices of cutification (*menghua*), state social media was also playing cute (*maimeng*), endeavoring to inspire affection by displaying its ability to engage in play.

The efficacy of state-run social media's adoption of a fandom register and interactive game modules such as the voting system beneath the livestream became evident on the nation's most public platforms. On the day People's Daily popularized the "Idols Group" hashtag, Baby Forklift became the number one "super topic" on Weibo list of fan circle trending topics. Pop idol fans themselves were shocked that a construction vehicle had eclipsed Mandopop singers and other celebrities (Figure 7). "I couldn't imagine the day would come that a forklift would occupy the number one position on the fan circle super topic list," one commented. Fan culture opinion leaders took note. A popular entertainment microblog posted, "Baby Forklift ranks first, followed by Lan Wangji [the blue digger] and Baby Mud Barfer. Construction Vehicle Fan Circle, come vote for the idols." This post takes a significant step toward resignifying remote supervisors as a fandom public, coining the name for a new fan circle (i挖掘机) according to the convention of adding "i," a homophone for "love" (爱, *ai*), to an idol's name—in this case, the Chinese words for excavating machine.

Figure 7. Ranking of celebrity idol topics trending in Weibo fan circles: Baby Forklift tops the list, followed by Mandopop singers Sean Xiao and KUN.

Enacting scripts of fandom—and inhabiting the role of fans—internet users responded to state social media's deployment of ludic modules of fandom management by further embellishing associations between the remote supervision and idol fandom in a number of creative ways. Posting names and images of the Idols Group to a dictionary of internet slang, one commented:

> *Produce Huoshenshan/Leishenshan Construction 101. [...] Remote supervisors, come and pick your favorite vehicles! [...Mine] is the hard-working and attractive Baby Forklift. Our Baby Forklift is the prettiest one in the field! Go for it, Baby Forklift! Wajijiwa, treat our brothers well!*

The post bristles with intertextual references to idol fan culture. "Produce Huoshenshan/Leishenshan Construction 101" is a clear reference to "Produce 101," a famous Chinese spinoff of a South Korean reality show that brings together trainees who compete for the opportunity to debut as an idols group. Wajijiwa (挖机机挖), which literally translates as "Dig machine dig," is the humorous nickname that fans gave to the company that managed the construction vehicles. It is a homophone for Wajijiwa (哇唧唧哇), a major Chinese entertainment conglomerate that manages numerous celebrity idols. This comment therefore imagines a fictional scenario with the Construction Vehicles Idols Group inhabiting the commercial spaces reserved for pop icons.

The collaborative resignification of the construction vehicle memes as celebrity idols was amplified by actual celebrities' engagement with the Idols Group hashtag, mobilizing further participation from their own circles of fans. For instance, Han Geng, a Chinese Mandopop singer and actor, sent a Weibo post using the Idols Group hashtag, stating, "Here comes a remote supervisor. We are supervising this significant national project together, and protecting the most mighty Baby Forklift together! Go for it, Wuhan! Go for it, Chinese people!" Han's comment received more than 100,000 reposts. It crystallized the symbolic reciprocity between vehicular idols who work on behalf of the nation and the fans who "protect" them through acts of devotion such as voting on ranking lists.

6 Devotional Artifacts as Participatory Propaganda

In this section, we explore the diversity of devotional artifacts created by members of the fandom public. To an extent, this profusion of creativity is a testimony to state-run social media's success in cultivating and harnessing citizens' devotional affect and investment in these idolized vehicles. In what follows, we focus in particular on Baby Forklift, by many measures the most popular and adored member of the Idols Group (Figure 8). While scholars have noted that participatory media may pose particular threats of "creative insurgency" (Kraidy 2017) to authoritarian regimes, we highlight how fandom governance offers Chinese

Figure 8. Weibo post of Baby Forklift cartoon with pink text that reads, "protect the cutest Little Forklift in the world."

state social media a tool for channeling creative expression into participatory propaganda (Asmolov 2019) without directly enforcing "interpretive domination" (Weller 1994).

A source of Baby Forklift's perceived cuteness (*mengdian*) was its relatively small size compared with the towering cranes and bulky cement mixers. This helps explain its special popularity and the particular ardor of its fans. For example, on 30 January, one devotee posted on Weibo:

> *Protect the cutest and most hardworking Baby Forklift in the world. Everybody work hard! Let our Baby Forklift debut and take center stage (C位出道)! Baby Forklift you can fly, your Fan Group [i叉] will follow you forever! Let's go, Huoshen and Leishen [Hospitals]! Let's go, Wuhan! Let's go, China!*

This comment displays many features indexical of Chinese fandom registers, including specialized cultural references to "debuting and taking center stage" and the identification of Baby Forklift's fan circle as i叉, which might be roughly translated as "Fork Lovers." Moreover, the post contains two speech acts characteristic of fandom: apostrophic words of encouragement directed to the idol and exhortations to fellow fans to engage in additional acts of devoted support and protection. In a symbolic loop, the imaginative construction of Baby Forklift's cute vulnerability subjectifies fans through their own vulnerability to its cuteness, summoning displays of concern and support that ramify "scalar intimacies" (Pritzker & Perrino 2020) outward in a chain of metonymic associations extending from the hospitals, to the affected city, to the nation itself.

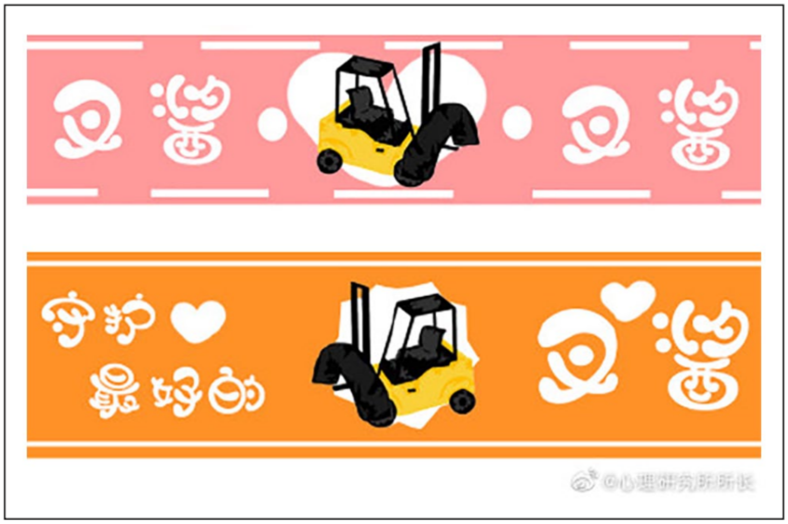

Figure 9. Baby Forklift fan banners.

In addition to producing and circulating memes and decorating photos as they might to show affection for celebrity entertainers, Fork Lovers used a variety of other expressive media to materialize affection for Baby Forklift and to augment intimacy with both the idol and other fans. For instance, they generated colorful banners emblazoned with Baby Forklift's image (Figure 9)—something that, in other circumstances, they might print and wave in the air at their idol's live appearances. Another fan created a video depicting Baby Forklift as a three-dimensional avatar inside an imaginary video game called "Let's Go Baby Forklift!" The video featured the plaintive sounds of a child singing, "heart to heart, we are a family," keying the affective bonds of kinship. In the accompanying text, the video's creator stated, "I am just a student, and don't have any money to donate. But I just want to contribute however I can." This comment reveals a key aspect of the devotion *meng* engenders: fans recognize the effort that an idol exerts, and reciprocate with strenuous displays of devotion—these acts are meaningful because, not in spite, of their insignificance. In this sense, vulnerability becomes the basis for agency (Pritzker 2020), realized through idols' devotion to duty, and fans' devotion to idols. Thus, fan art like these banners or videos served not only to generate emotional involvement with Baby Forklift and other idolized vehicles, but also to enact devotion as an ethical ideal and moral virtue.

Responsiveness in the context of social media goes hand-in-hand with expressiveness. According to Silvio (2019, p. 105), "the cute object induces a kind of mirroring effect, cutifying (and in some sense making abject) the viewing subject." In comments directed apostrophically to Baby Forklift, fans often used a Mandarin baby talk register (Farris 1992), infantilizing their speech as feminized Chinese caregivers might toward small children. Through the mirroring effect of baby talk, the cutification of the vehicles and cutification of the fans themselves converge in a single cute utterance: "Forky-Forky, charge forth for Mommy" (叉叉给妈妈冲).

Figure 10. Weibo comments glorifying Baby Forklift through creative combinations of visual and textual elements.

Adopting the stance of a mother figure toward the idol, as is common in female fan circles (Yan & Yang 2020), this commenter addresses it with a particularly childish diminutive.

In comment threads, fans graphically rendered effusive outpourings of emotion for Baby Forklift and other idols through paralinguistic features such as interjections, repetitions, punctuation marks (especially exclamation points), and *biaoqing* (表情)—emojis, emoticons, and stickers understood to express emotion (de Seta 2018) and pragmatically associated with playfulness (Y. Zhang et al. 2020). The most common emojis were hearts and cute animals, but Figure 10 shows a more elaborate example. In response to a People's Daily Weibo post about the vehicles, the first commenter embeds the phrases "top flow" (顶级流量)—meaning "top celebrity who generates the most data flow from online discourse"—and "cutey-cutey Little Baby Forklift" (可可爱爱小叉酱) in a long string of stars and moons (and one golden crown), iconizing the idol's position in the firmament. Decorative embellishment emphasizes devotional intensity.

The next commenter deploys parallelism (a common verbal feature in fan forums), echoing the alternation of key phrases with heavenly emojis, in this case, rainbows. The poetic comment reads:

> *Well-rounded top (idol) Little Forklift . . . Diligent Little Forklift . . . Pretty guy in winter, Little Forklift . . . Beautiful, kind-hearted Little Forklift . . . Consummately devoted to the country, Little Forklift . . . Born well and raised socialistically, Little Forklift . . . (If) you don't vote and I don't vote, when will Baby Forklift ascend?*

Like a magical incantation, this comment imbues the idol with enchantedness by rhythmically repeating its name (Malinowski 1965, p. 328). Together with the other examples in this section, from the banners to the video game, this illustrates fans' creative use of expressive forms to emotionally involve themselves with Baby Forklift.

7 Moral Panics Around Fandom Governance

Fans' digital displays of devotion created evidence of participation that amplified the intensity of the Idols Group sensation thereby spreading the reach of fandom governance. The heightened visibility of fandom practices around the vehicle memes also provided an object of discussion for internet users concerned about the COVID-19 crisis but less enamored with the vehicles themselves. In this final section, we examine some of the criticisms Chinese observers directed at Idols Group fandom. Mostly, critics focused on the dissonance between the seriousness of the unfolding crisis and fandom as a mode of creating an entertaining distraction. Many zeroed in on the frivolity of cuteness practices and their role in culturally constructing anthropomorphized objects of fan idolization. Taken together, these comments suggest a fragmentary reckoning with state social media's use of fandom governance as a mode of ludic governmentality.

From 29 to 30 January, the apex of the Idols Group's celebrity, observers on Chinese social media began to voice concerns about the unseemly intrusion of fandom during a national crisis. Among other things, critics worried that the significance of the idolized vehicles as metonymic extensions for the nation had been eclipsed by fetishistic adoration (as reflected, for instance, in the spell-like evocation of Baby Forklift in the previous example). "I disagree with the promotion of anthropomorphism," someone posted on Weibo:

> *but, surprisingly, I can accept the behavior of showing support to construction vehicles. I think we are actually showing support to workers sitting inside the vehicles. But don't go too far. No need to puff the list for the vehicles in [fan circles].*

Critics posted comments like these on the same platforms and in the same threads that Idols Group fans were using, casting doubts on the validity of emotionally inflamed vehicle fandom as a legitimate form of patriotic expression.

State social media accounts, which were deeply—and visibly—involved in animating the vehicles as idols, at times endeavored to realign fans' emotional expression with more explicitly patriotic concerns. For instance, in one Weibo post, the account of a CCTV subsidiary proclaimed, "the construction site is like a battlefield. Little Green, Big Yellow, Emperor Who Sends Dust, Emperor Who Lifts Things, and many other vehicles have joined the war. Let us pray together." Shifting to the metaphorical domain of warfare, this last sentence enjoins remote supervisors to pray (祈福), as women would have traditionally done for husbands, brothers, and sons off in battle. This injunction could therefore be read as an effort to realign the feminized idiom of fandom with a more traditionally pious female role, relegated to the domestic sphere. Responses to this post in the accompanying comment thread resisted this realignment, highlighting the fundamental intractability of play: "Where is my army of Baby Forklifts?" one commenter demanded. "Don't forget cute Baby Mud Barfer!" responded another. These irreverent responses point to the difficulty of constraining playfulness to the confines of ideologically delimited nationalist projects, and reveal the risk state social media runs of losing control of interpretations as it seeks to capitalize on the polyphony of fans' devotional behavior (Lei 2017).

Underlining the limits of fandom governance, a number of critics decried the way the Idols hullabaloo was distracting people from the controversy surrounding the Red Cross's mishandling of medical supplies. Some questioned ludic governmentality as a social media strategy, as in the following Weibo comment: "some things are not appropriate for entertainment . . . It is okay for internet users to joke around, but state-run media should not lead people in joking around." This comment makes clear that the state's role in animating the vehicles and stoking the flames of fandom was not lost on observers, who questioned mixing flippant playfulness with serious leadership in a time of crisis. Similarly, Huang Dianlin, a professor at the Communication University of China, posted an article titled "It is shameful for the media to cutify (*menghua*) the crisis" to his personal WeChat account. "The media is gradually losing . . . the care it should give the public," he wrote. "When everything is personified as cute stuff, there is no humanity. [. . .] The media should not treat crises as entertainment and games, and should not use frivolous and childish expressions to diminish the seriousness of social issues." Recall that some remote supervisors had framed state social media's playfulness as a form of parental care; here, Huang inverts this reading, suggesting that state social media's excessive playfulness and indulgence in *meng* reflects a dereliction of care. Before it was blocked, this post was read over 100,000 times.

Calling attention to the artificiality of treating vehicles like celebrities, some angry comments complained about the flagrant *inauthenticity* of the Idols sensation ("don't use your fake fan speech" or "stop mimicking fan speech in such an awkward way"), pointing to the kind of cynicism the state's social media influence campaigns risk engendering (Han 2015b). Nevertheless, numerous commentators seized upon the memeification and subsequent idolization of the vehicles as an occasion to articulate moral panics about "fan speech" and the stereotypical types of speakers they associated with it (Jones & Schieffelin 2009; Miller 2011). Critics construed emblematic features of fans' verbal and semiotic practices—nicknaming most salient among them—as inappropriate during a national crisis. For these critics, the excesses of *meng* was not just objectionable in the case of the Idols Group sensation, but reflective of deeper generational shortcomings (immaturity, frivolity, imbecility . . .) that the crisis only served to reveal.

For instance, the celebrity Daghe commented that people who think names like Baby Mud Barfer are "cute" (*meng*) display "indecent" (不得体) and "illogical" (没有逻辑) behavior. He called them "behind the times" (很土) and alleged that they have "low IQ" (智商不高) and "a very crude sense of cuteness" (萌点很土，总之不时尚). The allegation of "crude" *mengdian* has the dual connotation that the vehicles lack legitimate qualities of cuteness and that their fans have deficient judgment in discerning cuteness. Endorsing this interpretation, another user responded: "young people in our country, all of them are giant babies, silly kids, dummies. Giving names for vehicles, puffing ranking lists, and scoring points for idols. Calling this brother, that sister, this baby, that kissy-kissy." These comments connect a surfeit of *meng* and a childish desire for intimacy. Making this point more explicit, another commenter complained that nicknaming was "making the whole country childish," and singled out Little Corona as the most egregious example: "people are spoiling (宠溺) [the virus] with affection." In other words, the verbal incorporation of something deeply noxious into a discursive web of familial intimacy represented a childish overindulgence in cuteness.

The following Weibo comment from February 3 distills the essence of these criticisms into one overarching social critique:

> *I really don't like personifying everything. I think that it is embarrassing instead of cute. For instance, the forklift in the live video streaming a few days ago, and [the characters] Tender Country, Blazing Flag. Don't use fan culture for everything. I don't know why official Weibo accounts encourage this kind of weird list puffing and gawking. And Older Brother A-Zhong and Wuhan the Little Fool, I'm going to vomit. The thing I dislike is purposefully stirring up emotions. I want to slap the guy who does this. So disgusting.*

This comment offers a final reframing of the shifting modalities of visual participation that we have traced in this article. What began with Yangshipin's invitation for spectators to *witness*, and was then redefined in more participatory terms as remote *supervision* and parental *watching over*, is here negatively reevaluated as mere "*gawking*" (围观)—a term connoting the prurient gaze of idle onlookers, captivated by personified quasi-idols. This comment connects the animation of the Construction Vehicles Idols Group with several other parallel phenomena: the CCP's abortive effort to personify itself through anime characters Tender Country and Blazing Flag, the personification of the nation as Older Brother A-Zhong in fandom nationalism, and the proliferation of anthropomorphic nicknames like Little Fool surrounding COVID-19 crisis. It argues that these idol-esque personae arouse misplaced emotions, and questions why state social media would promote them as objects of fan-like engagement.

By generating cute personae that enlist citizen-fans as caring subjects, personification seems to be a vital component of fandom governance in the episode we analyze, as well as other events critical observers associated with it. By and in large, these critics did not reject fandom governance per se. Rather, they dwelled on the obtrusiveness of its emblematic cutified avatars and the diversion that they created from real problems. "Your country is called Brother A-Zhong," wrote one commenter, "the construction vehicle is called Baby Forklift, Wuhan City is called Little Fool. Everything can be personified. However, those workers, peasants, and minority groups who suffer a lot are not considered as real people." In a scathing critique of Chinese neoliberal consumerism, this commenter connected the personification of things according to the conventions of fandom with the dehumanizing treatment of actual persons — persons who toil, suffer, and die — as things. The key grievance about fandom, as both a cultural practice and a mode of governance, that surfaces here is that it distracts. Celebrity idols—be they human entertainers or apotheosized construction vehicles—have the potential to capture attention and transfix the gaze. Through ludic modules like ranking lists that cast idols as vulnerable, state social media harnessed this potential on a large scale; so much so that it prompted the observers whose comments we analyze here to critically reflect on the place of cuteness in Chinese public life more broadly.

8 Conclusion: The Cuteness of the State

Recent research on the politics of internet memes has shown that harnessing the creative potential of social media play can empower activists and energize social movements in dramatic ways, for instance the alt-right trolls who helped propel Donald Trump to the U.S. Presidency (Woods & Hahner 2019). Yet, An Xiao Mina (2019) argues, the power of memes to catalyze protest movements also makes them particularly appealing as tools of social influence for a Chinese state heavily invested in shaping social media dynamics. This paradox points to what Phillips and Milner (2017) describe as the "ambivalence" of the internet, characterized by the capacity of memes to both liberate and oppress, or of social media to both empower participation and constrain it. *Ludic governmentality* is our way of making sense of how Chinese state social media negotiated elements of playful co-production to harness the political potential of memes in a paternalistic form that blended care and control.

Moreover, as all these authors show, because of their open-ended meanings and almost infinite resignifiability, the same memes can be used for different purposes, political or otherwise. Drawing on fandom's register of verbal, visual, semiotic, and interactional practices, internet users "animated" (Silvio 2010, 2019) the construction vehicles from hospital construction livestreams, rendered them cute, and engaged with them in small but significant acts of symbolic devotion. By playfully embellishing an otherwise straightforward propagandistic spectacle, remote supervisors produced memetic content that Chinese state social media could itself appropriate and channel into a fandom apparatus it adapted from the entertainment industry. Although the Chinese state had suffered several previous failures in attempting to clothe itself in the mantle of cuteness, in this case, by responding to displays of cutification with culturally competent displays of playing cute, state social media accounts successfully harnessed the power of *meng* as a cultural matrix of devotional attachment. Embedding cute memes within modules of fandom management such as celebrity ranking lists, state social media co-produced the ephemeral sensation of the Construction Vehicle Idols group to implement what we have termed *fandom governance*.

Viewing Chinese state social media's co-production of Idols Group as an instance of fandom governance points to new capacities for engaging in playful—hence participatory—forms of ludic governmentality. In this article, we have highlighted how cuteness can be a source of political power—an ability to rally support through strategic displays of vulnerability. The construction vehicles could embody the cuteness of the state because state social media scaffolded playful engagement with them. Try as it might, the state cannot cutify itself. That requires the voluntary expressive labor of adoring fans, whose visible participation is a crucial part of successfully generating competitive devotion to cute idols. Yet the amount of public bandwidth taken up by these fandom practices in the midst of a pandemic also repelled some observers, who decried the excesses of cuteness and registered their disgust about misplaced intimacies. They were quick to note the performative nature of cuteness: there was nothing objectively cute about the vehicles or anything else associated with COVID-19 for that matter. Against the backdrop of the crisis, they alleged, the virality of the Idols Group was a perverse figment of playful fans cutifying everything and of the state brazenly playing cute.

These negative reactions point to some of the risks and limits of fandom governance. Recall W. Zhang's (2016) argument that visible fan participation can be a key part of an idol's allure. When Chinese state social media routed the construction vehicle memes into the Construction Vehicle Idols Group fandom, its tactics for amplifying fan engagement generated massive participation and publicity. Pitting cutified avatars against each other in ludic modules such as vote-fueled rankings, state social media capitalized on internet users' devotion to these vulnerable cuties. The resulting amplification of fans' emotional involvement fed into a large-scale spectacle of collective devotional behavior being visibly mobilized by state social media. But that spectacle was, in turn, available for Chinese social media users to scrutinize. A number of critical observers saw participation in Idols Group fandom as a sign not of social solidarity and patriotic devotion, but rather of social decay and unruly desire. Because fandom publics lay bare their devices of participation, fandom governance similarly exposes its own operation as a mechanism of social control.

One way of reading the Construction Vehicle Idols Group sensation is that state social media calculatedly embellished the vulnerability of the cute, but ultimately disposable, stand-ins for the state. In that sense, the cuteness of the state might appear to be an inconsequential part of the Chinese government's vast machinery for responding to the crisis and managing public perception of its response. Nevertheless, seeking recourse to the power of cuteness revealed the urgency for the state to make itself loveable at a time of national crisis. When it plays cute, the state reveals a dependence on subjects' devotion akin to idols' need for their fans. By playing the devotional games of fandom governance, remote supervisors were conscripted into protecting not only the manufactured vulnerability of Baby Forklift and the other Idols, but also the symbolic vulnerability of the state itself. Cuteness is power.

References for Chapter 1

Abidin, Crystal. "Agentic Cute (^.^): Pastiching East Asian Cute in Influencer Commerce." *East Asian Journal of Popular Culture* 2, no. 1 (April 1, 2016): 33–47. https://doi.org/10.1386/eapc.2.1.33_1.

Agha, Asif. "Register." *Journal of Linguistic Anthropology* 9, no. 1/2 (1999): 216–19.

Ang, Yuen Yuen. "When COVID-19 Meets Centralized, Personalized Power." *Nature Human Behaviour* 4, no. 5 (May 2020): 445–47. https://doi.org/10.1038/s41562-020-0872-3.

Asmolov, Gregory. "The Effects of Participatory Propaganda: From Socialization to Internalization of Conflicts." *Journal of Design and Science*, no. 6 (August 7, 2019). https://doi.org/10.21428/7808da6b.833c9940.

Bateson, Gregory. *Steps to an Ecology of Mind: Collected Essays in Anthropology, Psychiatry, Evolution, and Epistemology*. Chicago, IL: University of Chicago Press, 2000. https://press.uchicago.edu/ucp/books/book/chicago/S/bo3620295.html.

Bonilla, Yarimar, and Jonathan Rosa. "#Ferguson: Digital Protest, Hashtag Ethnography, and the Racial Politics of Social Media in the United States." *American Ethnologist* 42, no. 1 (2015): 4–17. https://doi.org/10.1111/amet.12112.

Brady, Anne-Marie. "Regimenting the Public Mind: The Modernization of Propaganda in the PRC." *International Journal* 57, no. 4 (December 1, 2002): 563–78. https://doi.org/10.1177/002070200205700404.

Cao, Xuenan. "Bullet Screens (Danmu): Texting, Online Streaming, and the Spectacle of Social Inequality on Chinese Social Networks." *Theory, Culture & Society* 38, no. 3 (May 2021): 29–49. https://doi.org/10.1177/0263276419877675.

Coates, Jamie. "So 'Hot' Right Now Reflections on Virality and Sociality from Transnational Digital China." *Digital Culture & Society* 3, no. 2 (December 20, 2017): 77–98. https://doi.org/10.14361/dcs-2017-0206.

De Seta, Gabriele. "'Meng? It Just Means Cute': A Chinese Online Vernacular Term in Context." *M/C Journal* 17, no. 2 (March 3, 2014). https://doi.org/10.5204/mcj.789.

Douban. 领风者 *[The Leader]*, 2019. https://movie.douban.com/subject/30409395/.

Fang, Kecheng, and Maria Repnikova. "Demystifying 'Little Pink': The Creation and Evolution of a Gendered Label for Nationalistic Activists in China." *New Media & Society* 20, no. 6 (June 1, 2018): 2162–85. https://doi.org/10.1177/1461444817731923.

Farris, Catherine S. "Chinese Preschool Codeswitching: Mandarin Babytalk and the Voice of

Authority." *Journal of Multilingual and Multicultural Development* 13, no. 1–2 (January 1, 1992): 187–213. https://doi.org/10.1080/01434632.1992.9994491.

Feng, Xiaodan. "Curating and Exhibiting for the Pandemic: Participatory Virtual Art Practices During the COVID-19 Outbreak in China." *Social Media + Society* 6, no. 3 (July 1, 2020): 2056305120948232. https://doi.org/10.1177/2056305120948232.

Foucault, Michel. *The Foucault Effect: Studies in Governmentality: With Two Lectures by and an Interview with Michel Foucault*. Edited by Graham Burchell, Colin Gordon, and Peter Miller. Chicago: University of Chicago Press, 1991.

Goffman, Erving. *Frame Analysis: An Essay on the Organization of Experience*. Harvard University Press, 1974.

Gorwa, Robert. "What Is Platform Governance?" *Information, Communication & Society* 22, no. 6 (May 12, 2019): 854–71. https://doi.org/10.1080/1369118X.2019.1573914.

Guo, Ting. "Politics of Love: Love as a Religious and Political Discourse in Modern China through the Lens of Political Leaders." *Critical Research on Religion* 8, no. 1 (April 1, 2020): 39–52. https://doi.org/10.1177/2050303219874366.

Han, Rongbin. "Defending the Authoritarian Regime Online: China's 'Voluntary Fifty-Cent Army.'" *The China Quarterly*, no. 224 (2015): 1006–25.

——. "Manufacturing Consent in Cyberspace: China's 'Fifty-Cent Army.'" *Journal of Current Chinese Affairs* 44, no. 2 (June 1, 2015): 105–34. https://doi.org/10.1177/186810261504400205.

Hiramoto, Mie, and Lionel Wee. "*Kawaii* in the Semiotic Landscape." *Sociolinguistic Studies* 13, no. 1 (April 13, 2019): 15–35. https://doi.org/10.1558/sols.36212.

Initium Media. "阿中哥哥只有我們了: 國家機構支持出征海外社交平台，你如何看飯圈式愛國? [Brother A-Zhong Only Has Us: State institutions supporting expeditions to overseas social media, what do you think of fan- style patriotism?]." 端傳媒 Initium Media, August 16, 2019. https://theinitium.com/roundtable/20190816-roundtable-zh-a-zhong/.

Johnson, Amy. "Decrowning Doubles: Indexicality and Aspect in a Bahraini Twitter Parody Account." *Al-'Arabiyya* 48 (2015): 61–83.

Jones, Graham M., and Bambi B. Schieffelin. "Talking Text and Talking Back: 'My Bff Jill' from Boob Tube to YouTube." *Journal of Computer-Mediated Communication* 14, no. 4 (July 1, 2009): 1050–79. https://doi.org/10.1111/j.1083-6101.2009.01481.x.

Jung, Sun. *Korean Masculinities and Transcultural Consumption: Yonsama, Rain, Oldboy, K-Pop Idols*. Hong Kong University Press, 2011.

Kelty, Christopher M. *The Participant: A Century of Participation in Four Stories*. Chicago, IL: University of Chicago Press, 2020. https://press.uchicago.edu/ucp/books/book/chicago/P/bo44520895.html.

Kraidy, Marwan. “Creative Insurgency and the Celebrity President: Politics and Popular Culture from the Arab Spring to the White House.” *Arab Media and Society*, no. 23 (Winter/Spring 2017). https://www.arabmediasociety.com/wp-content/uploads/2017/12/20170118124014_Creative_Insurgency.pdf.

Lagerkvist, Johan. “Internet Ideotainment in the PRC: National Responses to Cultural Globalization.” *Journal of Contemporary China* 17, no. 54 (February 1, 2008): 121–40. https://doi.org/10.1080/10670560701693120.

Lei, Ya-Wen. *The Contentious Public Sphere*. Princeton: Princeton University Press, 2017. https://press.princeton.edu/books/hardcover/9780691166865/the-contentious-public-sphere.

Liu, Hailong. *From Cyber-Nationalism to Fandom Nationalism: The Case of Diba Expedition in China*. London; New York: Routledge, Taylor & Francis Group, 2020.

Malinowski, Bronislaw. *Coral Gardens and Their Magic, Volume II: The Language of Magic and Gardening*. 1st edition. Indiana University Press, 1965.

Miller, Laura. “Subversive Script and Novel Graphs in Japanese Girls’ Culture.” *Language & Communication* 31, no. 1 (January 2011): 16–26. https://doi.org/10.1016/j.langcom.2010.11.003.

Miltner, Kate M. “Internet Memes.” In *The SAGE Handbook of Social Media*, edited by Jean Burgess, Alice Marwick, and Thomas Poell. London: SAGE Publications Ltd, 2018. https://doi.org/10.4135/9781473984066.

Mina, An Xiao. *Memes to Movements: How the World’s Most Viral Media Is Changing Social Protest and Power*. Boston: Beacon Press, 2019.

Moore, Robert L. “Nicknames in Urban China: A Two-Tiered Model.” *Names* 41, no. 2 (June 1, 1993): 67–86. https://doi.org/10.1179/nam.1993.41.2.67.

Moskowitz, Marc L. *Internet Video Culture in China: YouTube, Youku, and the Space in Between*. First published 2019. First issued in paperback 2021. Routledge Contemporary China Series. Routledge, 2021.

Ngai, Sianne. *Our Aesthetic Categories: Zany, Cute, Interesting*. Cambridge, Mass.; London: Harvard University Press, 2015.

Palmer, David A. *Qigong Fever: Body, Science, and Utopia in China*. Columbia University Press, 2007.

Phillips, Whitney, and Ryan M. Milner. *The Ambivalent Internet: Mischief, Oddity, and Antagonism Online*. 1st edition. Cambridge, UK ; Malden, MA: Polity, 2017.

Phoenix Media. "江山娇与红旗漫"为何仅5 小时就被网友骂到下架 [Why Tender Country and Blazing Flag only survived for five hours until netizen criticisms forced them out].," February 20, 2020. https://ishare.ifeng.com/c/s/7uEsfN3EhPu.

Pink, Sarah, Heather A. Horst, John Postill, Larissa Hjorth, Tania Lewis, and Jo Tacchi, eds. *Digital Ethnography: Principles and Practice*. Los Angeles: SAGE, 2016.

Pritzker, Sonya E. "Language, Emotion, and the Politics of Vulnerability." *Annual Review of Anthropology* 49, no. 1 (October 21, 2020): 241–56. https://doi.org/10.1146/annurev-anthro-010220-074429.

Pritzker, Sonya E., and Sabina Perrino. "Culture inside: Scale, Intimacy, and Chronotopic Stance in Situated Narratives." *Language in Society* 50, no. 3 (June 2021): 365–87. https://doi.org/10.1017/S0047404520000342.

Qin, Amy. "China Pledged to Build a New Hospital in 10 Days. It's Close." *The New York Times*, February 3, 2020, sec. World. https://www.nytimes.com/2020/02/03/world/asia/coronavirus-wuhan-hospital.html.

Rea, Christopher G. "Spoofing (e'gao) Culture on the Chinese Internet." In *Humour in Chinese Life and Culture*. Hong Kong University Press, 2013. https://doi.org/10.5790/hongkong/9789888139231.003.0007.

Repnikova, Maria, and Kecheng Fang. "Authoritarian Participatory Persuasion 2.0: Netizens as Thought Work Collaborators in China." *Journal of Contemporary China* 27, no. 113 (September 3, 2018): 763–79. https://doi.org/10.1080/10670564.2018.1458063.

Seta, Gabriele de. "Biaoqing: The Circulation of Emoticons, Emoji, Stickers, and Custom Images on Chinese Digital Media Platforms." *First Monday*, September 1, 2018. https://doi.org/10.5210/fm.v23i9.9391.

Chapter 2

Viral Visualizations: How Coronavirus Skeptics Use Orthodox Data Practices to Promote Unorthodox Science Online

Abstract

Controversial understandings of the coronavirus pandemic have turned data visualizations into a battleground. Defying public health officials, coronavirus skeptics on US social media spent much of 2020 creating data visualizations showing that the government's pandemic response was excessive and that the crisis was over. This paper investigates how pandemic visualizations circulated on social media, and shows that people who mistrust the scientific establishment often deploy the same rhetoric of data-driven decision-making used by experts, but to advocate for radical policy changes. Using a quantitative analysis of how visualizations spread on Twitter and an ethnographic approach to analyzing conversations about COVID data on Facebook, we document an epistemological gap that leads pro- and anti-mask groups to draw drastically different inferences from similar data. Ultimately, we argue that the deployment of COVID data visualizations reflect a deeper sociopolitical rift regarding the place of science in public life.

Publication: This chapter appears as Crystal Lee, Tanya Yang, Gabrielle Inchoco, Graham M. Jones, and Arvind Satyanarayan, Honorable "Viral Visualizations: How Coronavirus Skeptics Use Orthodox Data Practices to Promote Mention Unorthodox Science Online." *ACM Human Factors in Computing Systems (CHI)*, 2021.

Acknowledgements: The authors thank Stephan Risi, Maeva Fincker, and Mateo Mon- terde for their assistance with quantitative methods and supplemental material. We also thank the members of the Visualization Group (especially Jonathan Zong, Alan Lundgard, Harini Suresh, EJ Sefah, and the Fall 2020 UROP cohort: Anna Arpaci-Dusseau, Anna Meurer, Ethan Nevidomsky, Kat Huang, and Soomin Chun). This manuscript benefited from the insights of Rodrigo Ochigame, Hannah LeBlanc, Meghan Kelly, Will Deringer, Blakeley H. Payne, Mariel García-Montes, and the comments of anonymous reviewers. This project was supported by NSF Award 1900991, NSF Dissertation Improvement Grant 1941577, an SSRC Social Data Dissertation Fellowship, and the MIT Programs for Digital Humanities.

1 Introduction

Throughout the coronavirus pandemic, researchers have held up the crisis as a “breakthrough moment” for data visualization research (Shneiderman 2020): John Burn-Murdoch's line chart comparing infection rates across countries helped millions of people make sense of the pandemic's scale in the United States (Forrest 2020), and even top Trump administration officials seemed to rely heavily on the Johns Hopkins University COVID data dashboard (Mazza 2020). Almost every US state now hosts a data dashboard on their health department website to show how the pandemic is unfolding. However, despite a preponderance of evidence that masks are crucial to reducing viral transmission (CDC 2020, Chu et al. 2020, Su et al. 2020), protestors across the United States have argued for local governments to overturn their mask mandates and begin reopening schools and businesses. A pandemic that affects a few, they reason, should not impinge on the liberties of a majority to go about life as usual. To support their arguments, these protestors and activists have created thousands of their own visualizations, often using the same datasets as health officials.

This paper investigates how these activist networks use rhetorics of scientific rigor to oppose these public health measures. Far from ignoring scientific evidence to argue for individual freedom, anti-maskers often engage deeply with public datasets and make what we call “**counter-visualizations**”—visualizations using orthodox methods to make unorthodox arguments—to challenge mainstream narratives that the pandemic is urgent and ongoing. By asking community members to “follow the data,” these groups mobilize data visualizations to support significant local changes.

We examine the circulation of COVID-related data visualizations through both quantitative and qualitative methods. First, we conduct a quantitative analysis of close to half a million tweets that use data visualizations to talk about the pandemic. We use network analysis to identify communities of users who retweet the same content or otherwise engage with one another (e.g., maskers and anti-maskers). We process over 41,000 images through a computer vision model trained by Poco and Heer (2017), and extract feature embeddings to identify clusters and patterns in visualization designs. The academic visualization research community has traditionally focused on mitigating chartjunk and creating more intuitive visualization tools for use by non-experts; *better* visualizations, researchers argue, would aid public understanding of data-driven phenomena. However, we find that anti-mask groups on Twitter often create polished counter-visualizations that would not be out of place in scientific papers, health department reports, and publications like the *Financial Times*.

Second, we supplement this quantitative work with a six month-long observational study of anti-mask groups on Facebook. The period of this study, March to September 2020, was critical as it spanned the formation and consolidation of these groups at the pandemic's start. Quantitative analysis gives us an overview of what online discourse about data and its visual representation looks like on Twitter both within and outside anti-mask communities. Qualitative analysis of anti-mask groups gives us an interactional view of how these groups leverage the language of scientific rigor—being critical about data sources, explicitly stating analytical limitations of specific models, and more—in order to support ending public health restrictions despite the consensus of the scientific establishment. Our data analysis evolved as these communities did, and our methods reflect how these users reacted in real time to the kaleidoscopic nature of pandemic life. As of this writing, Facebook has banned some of the groups we studied, who have since moved to more unregulated platforms (Parler and MeWe).

While previous literature in visualization and science communication has emphasized the need for data and media literacy as a way to combat misinformation (Fontichiaro and Oehrli 2016, Guess et al. 2020, Scheufele and Krause 2019), this study finds that anti-mask groups practice a form of data literacy in spades. Within this constituency, unorthodox viewpoints do not result from a deficiency of data literacy; sophisticated practices of data literacy are a means of consolidating and promulgating views that fly in the face of scientific orthodoxy. Not only are these groups prolific in their creation of counter-visualizations, but they leverage data and their visual representations to advocate for and enact policy changes on the city, county, and state levels.

As we shall see throughout this case study, anti-mask communities on social media have defined themselves *in opposition* to the discursive and interpretive norms of the mainstream public sphere (e.g., against the “lamestream media”). In media studies, the term “counterpublic” describes constituencies that organize themselves in opposition to mainstream civic discourse, often by agentively using communications media (Downey and Fenton 2003). In approaching anti-maskers as a counterpublic (a group shaped by its hostile stance toward mainstream science), we focus particular attention on one form of agentive media production central to their movement: data visualization. We define this counterpublic's visualization practices as “counter-visualizations” that use orthodox scientific methods to make unorthodox arguments, beyond the pale of the scientific establishment. Data visualizations are not a neutral window onto an observer-independent reality; during a pandemic, they are an arena of political struggle.

Among other initiatives, these groups argue for open access to government data (claiming that CDC and local health departments are not releasing enough data for citizens to make informed decisions), and they use the language of data-driven decision-making to show that social distancing mandates are both ill-advised and unnecessary. In these discussions, we find that anti-maskers think carefully about the grammar of graphics by decomposing visualizations into layered components (e.g., raw data, statistical transformations, mappings, marks, colors). Additionally, they debate how each component changes the narrative that the visualization tells, and they brainstorm alternate visualizations that would better enhance public understanding of the data. This paper empirically shows how COVID anti-mask groups use data visualizations to argue that the US government's response (broadly construed) to the pandemic is overblown, and that the crisis has long been over.

These findings suggest that the ability for the scientific community and public health departments to better convey the urgency of the US coronavirus pandemic may not be strengthened by introducing more downloadable datasets, by producing “better visualizations” (e.g., graphics that are more intuitive or efficient), or by educating people on how to better interpret them. This study shows that there is a fundamental epistemological conflict between maskers and anti-maskers, who use the same data but come to such different conclusions. As science and technology studies (STS) scholars have shown, data is not a neutral substrate that can be used for good or for ill (Bowker and Star 2000, Gitelman 2013, Porter 1995). Indeed, anti-maskers often reveal themselves to be more sophisticated in their understanding of how scientific knowledge is socially constructed than their ideological adversaries, who espouse naive realism about the “objective” truth of public health data. Quantitative data is culturally and historically situated; the manner in which it is collected, analyzed, and interpreted reflects a deeper narrative that is bolstered by the collective effervescence found within social media communities. Put differently, there is no such thing as dispassionate or objective data analysis. Instead, there are stories: stories shaped by cultural logics, animated by personal experience, and entrenched by collective action. This story is about how a public health crisis—refracted through seemingly objective numbers and data visualizations—is part of a broader battleground about scientific epistemology and democracy in modern American life.

2 Related Work

2.1 Data and visualization literacies

There is a robust literature in computer science on data and visualization literacy, where the latter often refers to the ability of a person “to comprehend and interpret graphs” (Lee, Kim, and Kwon 2017) as well as the ability to create visualizations from scratch (Börner, Bueckle, and Ginda 2019). Research in this area often includes devising methods to assess this form of literacy (Alper et al. 2017, Boy et al. 2014, Lengler 2006), to investigate how people (mis)understand visualizations (Börner et al. 2016; Valdez, Ziefle, and Sedlmair 2018; Parsons 2018), or to create systems that help a user improve their understanding of unfamiliar visualizations (Adar and Lee 2020; Alberda 2020; Börner, Bueckle, and Ginda 2019; Ruchikachorn and Mueller 2015; Stoiber et al. 2019; Tanahashi, Leaf, and Ma 2016). Evan Peck et al. (2019) have responded to this literature by showing how a *“complex tapestry of motivations, preferences, and beliefs* [impact]*the way that participants* [prioritize]*data visualizations,”* suggesting that researchers need to better understand users’ social and political context in order to design visualizations that speak powerfully to their personal experience.

Linguistic anthropologists have shown that "literacy" is not just the ability to encode and decode written messages. The skills related to reading and writing are historically embedded, and take on different meanings in a given social context depending on who has access to them, and what people think they should and shouldn't be used for. As a consequence of such contingencies, these scholars view literacy as multiple rather than singular, attending to the impact of local circumstances on the way members of any given community practice literacy and construe its value (Ahearn 2004). Thus, media literacy is not simply about understanding information, but about being able to actively leverage it in locally relevant social interactions (Hobbs 2016). Building on these traditions, we do not normatively assess anti-maskers' visualization practices against a prescriptivist model of what literacy *should* be (according to, say, experts in human-computer interaction), but rather seek to describe what those practices actually look like in locally relevant contexts.

As David Buckingham (2017) has noted, calls for increased literacy have often become a form of wrong-headed solutionism that posits education as the fix to all media-related problems. danah boyd (2018) has documented, too, that calling for increased media literacy can often backfire: the instruction to "question more" can lead to a weaponization of critical thinking and increased distrust of media and government institutions. She argues that calls for media literacy can often frame problems like fake news as ones of personal responsibility rather than a crisis of collective action. Similarly, Francesca Tripodi (2018) has shown how evangelical voters do not vote for Trump because they have been "fooled" by fake news, but because they privilege the personal study of primary sources and have found logical inconsistencies not in Trump's words, but in mainstream media portrayals of the president. As such, Tripodi argues, media literacy is not a means of fighting "alternative facts." Christopher Bail et al. (2018) have also shown how being exposed to more opposing views can actually increase political polarization.

Finally, in his study of how climate skeptics interpret scientific studies, Frank Fischer (2019) argues that increasing fact-checking or levels of scientific literacy is insufficient for fighting alternative facts. "While fact checking is a worthy activity," he says, "we need to look deeper into this phenomenon to find out what it is about, what is behind it." Further qualitative studies that investigate how ideas are culturally and historically situated, as the discussion around COVID datasets and visualizations are manifestations of deeper political questions about the role of science in public life.

2.2 Critical approaches to visualization

Historians, anthropologists, and geographers have long shown how visualizations—far from an objective representation of knowledge—are often, in fact, representations of power (Harley 1989; Kitchin, Dodge, and Perkins 2011; Mueggler 2011; Teng 2004). To address this in practice, feminist cartographers have developed quantitative GIS methods to describe and analyze differences across race, gender, class, and space, and these insights are then used to inform policymaking and political advocacy (Hanson 1992, Kwan 2002, McLafferty 1995). Best practices in data visualization have often emphasized reflexivity as a way to counter the power dynamics and systemic inequalities that are often inscribed in data science and design (Correll 2019, Costanza-Chock 2020, D'Ignazio and Klein 2020). Central to this practice is articulating what is excluded from the data (Buolamwini and Gebru 2018, Noble 2018, Onuoha 2016), understanding how data reflect situated knowledges rather than objective truth (Christin 2016, Haraway 1988, So and Duarte 2020), and creating alternative methods of analyzing and presenting data based on anti-oppressive practices (Anti-Eviction Mapping Project 2019, Data for Black Lives 2020, Kelly 2019). Researchers have also shown how interpreting data visualizations is a fundamentally social, narrative-driven endeavor (Dumit and Burri 2008; Hullman and Diakopoulos 2011; Peck, Ayuso, and El-Etr 2019). By focusing on a user's contextual experience and the communicative dimensions of a visualization, computer scientists have destabilized a more traditional focus on improving the technical components of data visualization towards understanding how users interpret and use them (Lee at al. 2016; Maltese, Svetina, and Harsh 2015, Viegas and Wattenberg 2006).

Critical, reflexive studies of data visualization are undoubtedly crucial for dismantling computational systems that exacerbate existing social inequalities. Research on COVID visualizations is already underway: Emily Bowe et al. (2020) have shown how visualizations reflect the unfolding of the pandemic at different scales; Alexander Campolo (2020) has documented how the pandemic has produced new forms of visual knowledge, and Yixuan Zhang et al. (2020) have mapped the landscape of COVID-related crisis visualizations. This paper builds on these approaches to investigate the epistemological crisis that leads some people to conclude that mask-wearing is a crucial public health measure and for others to reject it completely.

Like data feminists, anti-mask groups similarly identify problems of political power within datasets that are released (or otherwise withheld) by the US government. Indeed, they contend that the way COVID data is currently being collected is non-neutral, and they seek liberation from what they see as an increasingly authoritarian state that weaponizes science to exacerbate persistent and asymmetric power relations. This paper shows that more critical approaches to visualization are necessary, and that the frameworks used by these researchers (e.g., critical race theory, gender analysis, and social studies of science) are crucial to disentangling how anti-mask groups mobilize visualizations politically to achieve powerful and often horrifying ends.

3 Methods

This paper pairs a quantitative approach to analyzing Twitter data (computer vision and network analysis) with a qualitative approach to examining Facebook Groups (digital ethnography). Drawing from social media scholarship that uses mixed-methods approaches to examine how users interact with one another (Arif, Stewart, and Starbird 2018; Berriche and Altay 2020; Burgess and Matamoros-Fernández 2016; Kiesling et al. 2018; Moats and Borra 2018), this paper engages with work critical of quantitative social media methods (Baym 2013, Wu and Taneja 2020) by demonstrating how interpretive analyses of social media discussions and computational techniques can be mutually reinforcing. In particular, we leverage quantitative studies of social media that use network analysis to understand political polarization (Arif, Stewart, and Starbird 2018), qualitative analysis of comments to identify changes in online dialogue over time (Yardi and boyd 2010), and visualization research that reverse-engineers and classifies chart images (Savva et al. 2011; Poco and Heer 2017).

3.1 Twitter data and quantitative analysis

3.1.1 Dataset. This analysis is conducted using a dataset of tweet IDs that Emily Chen et al. (2020) assembled by monitoring the Twitter streaming API for certain keywords associated with COVID-19 (e.g., "coronavirus", "pandemic", "lockdown", etc.) as well as by following the accounts of public health institutions (e.g., @CDCGov, @WHO, etc). We used version 2.7 of this dataset, which included over 390M tweets spanning January 21, 2020–July 31, 2020. This dataset consists of tweet IDs which we "hydrated" using twarc (Documenting the Now 2016) into full tweets with associated metadata (e.g., hashtags, mentions, replies, favorites, retweets, etc.).

To identify tweets that primarily discuss data visualizations about the pandemic, we initially adopted a strategy that filtered for tweets that contained at least one image and keyword associated with data analysis (e.g., "bar," "line," but also "trend," "data"). Unfortunately, this strategy yielded more noise than signal as most images in the resultant dataset were memes and photographs. We therefore adopted a more conservative approach, filtering for tweets that explicitly mentioned chart-related keywords (i.e., "chart(s)", "plot(s)", "map(s)", "dashboard(s)", "vis", "viz", or "visualization(s)"). This process yielded a dataset of almost 500,000 tweets that incorporated over 41,000 images. We loaded the tweets and their associated metadata into a SQLite database, and the images were downloaded and stored on the file system.

3.1.2 Image classification. To analyze the types of visualization found in the dataset, we began by classifying every image in our corpus using the mark classification computer vision model trained by Poco and Heer (2017). Unfortunately, this model was only able to classify 30% of the images. As a result, we extracted a 4096-dimensional feature embedding for every image, and ran k-means clustering on a 100-dimensionally reduced space of these embeddings, for steps of k from 5–40. Two authors manually inspected the outputs of these runs, independently identified the most salient clusters, and then cross validated their analysis to assemble a final list of 8 relevant clusters: line charts, area charts, bar charts, pie charts, tables, maps, dashboards, and images. For dimensionality reduction and visualization, we used the UMAP algorithm (McInnes, Healy, and Melville 2018) and iteratively arrived at the following parameter settings: 20 neighbors, a minimum distance of 0.01, and using the cosine distance metric. To account for UMAP's stochasticity, we executed 10 runs and qualitatively examined the output to ensure our analyses were not based on any resultant artifacts.

3.1.3 Network analysis. Finally, to analyze the users participating in these discussions, we constructed a network graph: nodes were users who appeared in our dataset, and edges were drawn between pairs of nodes if a user mentioned, replied to, retweeted, or quote-tweeted another user. The resultant graph consisted of almost 400,000 nodes and over 583,000 edges, with an average degree of 2.9. To produce a denser network structure, we calculated a histogram of node degrees and identified that two-thirds of the nodes were degree 1. We then computed subgraphs, filtering for nodes with a minimum degree of 2 through 10 and found that degree 5 offered us a good balance between focusing on the most influential actors in the network, without losing smaller yet salient communities. This step yielded a subgraph of over 28,000 nodes and 104,000 edges, with an average degree of 7.3. We detected communities on this network using the Louvain method (Blondel et al. 2008). Of the 2,573 different communities detected by this algorithm, we primarily focus on the top 10 largest communities which account for 72% of nodes, 80% of edges, and 30% of visualizations.

3.2 Facebook data and qualitative analysis

3.2.1 Digital ethnography. While qualitative research can involve clinical protocols like interviews or surveys, Clifford Geertz (1998) argues that the most substantial ethnographic insights into the cultural life of a community come from “**deep hanging out**,” i.e., long-term, participant observation alongside its members. Using “lurking,” a mode of participation by observing specific to digital platforms, we propose “**deep lurking**” as a way of systematically documenting the cultural practices of online communities. Our methods here rely on robust methodological literature in digital ethnography (Coleman 2014, Markham 2013), and we employ a case study approach (Small 2009) to analyze these Facebook groups. To that end, we followed five Facebook groups (each with a wide range of followers, 10K-300K) over the first six months of the coronavirus pandemic, and we collected posts throughout the platform that included terms for “coronavirus” and “visualization” with Facebook's CrowdTangle tool (CrowdTangle Team 2020). In our deep lurking, we archived web pages and took field notes on the following: posts (regardless of whether or not they included “coronavirus” and “data”), subsequent comments, Facebook Live streams, and photos of in-person events. We collected and analyzed posts from these groups from their earliest date to September 2020.

Taking a case study approach to the interactional Facebook data yields an analysis that ultimately complements the quantitative analysis. While the objective with analyzing Twitter data is statistical representativeness—we investigate which visualizations are the most popular, and in which communities—the objective of analyzing granular Facebook data is to accurately understand social dynamics within a singular community (Small 2009). As such, the Twitter and Facebook analyses are foils of one another: we have the ability to quantitatively analyze large-scale interactions on Twitter, whereas we analyze the Facebook data by close reading and attending to specific context. Twitter communities are loosely formed by users retweeting, liking, or mentioning one another; Facebook groups create clearly bounded relationships between specific communities. By matching the affordances of each data source with the most ecologically appropriate method (network analysis and digital ethnography), this paper meaningfully combines qualitative and quantitative methods to understand data visualizations about the pandemic on a deeply contextual level and at scale.

3.2.2 Data collection & analysis. Concretely, we printed out posts as PDFs, tagged them with qualitative analysis software, and synthesized themes across these comments using grounded theory (Kozinets 2019). Grounded theory is an inductive method where researchers collect data and tag it by identifying analytically pertinent themes. Researchers then group these codes into higher-level concepts. As Kathy Charmaz (2006) writes: these "methods consist of systematic, yet flexible guidelines for collecting and analyzing qualitative data to construct theories 'grounded' in the data themselves," and these methods have since been adapted for social media analysis (Postill and Pink 2012). While this flexibility allows this method to respond dynamically to changing empirical phenomena, it can also lead to ambiguity about how new data fit with previously identified patterns. Digital ethnography also requires a longer time horizon than quantitative work in order to generate meaningful insights and, on its own, does not lead immediately to quantifiable results. These limitations are a major reason to use both qualitative and quantitative approaches. Following Emerson et al. (2011), we employ an integrative strategy that weaves together "exemplars" from qualitative data alongside our interpretations. We have redacted the names of individual users and the Facebook groups we have studied, but we have preserved the dates and other metadata of each post within the article where possible.

3.2.3 Note on terminology. Throughout this study, we use the term "anti-mask" as a synecdoche for a broad spectrum of beliefs: that the pandemic is exaggerated, schools should be reopening, etc. While groups who hold these beliefs are certainly heterogeneous, the mask is a common flashpoint throughout the ethnographic data, and they use the term "maskers" to describe people who are driven by fear. They are "anti-mask" by juxtaposition. This study therefore takes an emic (i.e. "insider") approach to analyzing how members of these groups think, talk, and interact with one another, which starts by using terms that these community members would use to describe themselves. There is a temptation in studies of this nature to describe these groups as "anti-science," but this would make it completely impossible for us to meaningfully investigate this article's central question: understanding what these groups mean when they say "science."

4 Case Study

In the Twitter analysis, we quantitatively examine a corpus of tweets that use data visualizations to discuss the pandemic, and we create a UMAP visualization (figure 1) that identifies the types of visualizations that proliferate on Twitter. Then, we create a network graph (figure 2) of the users who share and interact with these data visualizations; the edges that link users in a network together are retweets, likes, mentions. We discover that the fourth largest network in our data consists of users promulgating heterodox scientific positions about the pandemic (i.e., anti-maskers). By comparing the visualizations shared within anti-mask and mainstream networks, we discover that there is no significant difference in the kinds of visualizations that the communities on Twitter are using to make drastically different arguments about coronavirus (figure 3). Anti-maskers (the community with the highest percentage of verified users) also share the second-highest number of charts across the top six communities (table 1), are the most prolific producers of area/line charts, and share the fewest number of photos (memes and images of politicians; see figure 3). Anti-maskers are also the most likely to amplify messages from their own community. We then examine the kinds of visualizations that anti-maskers discuss (figure 4).

This leads us to an interpretive question that animates the Facebook analysis: how can opposing groups of people use similar methods of visualization and reach such different interpretations of the data? We approach this problem by ethnographically studying interactions within a community of anti-maskers on Facebook to better understand their practices of knowledge-making and data analysis, and we show how these discussions exemplify a fundamental epistemological rift about how knowledge about the coronavirus pandemic should be made, interpreted, and shared.

4.1 Visualization design and network analysis

4.1.1 Visualization types. What kinds of visualizations are Twitter users sharing about the pandemic? Figure 1 visualizes the feature embeddings of images in our corpus, with color encoding clusters revealed and manually curated through k-means. Each circle is sized by the engagement the associated tweet received calculated as the sum of the number of favorites, replies, retweets, and quote tweets. Our analysis revealed eight major clusters: line charts (8908 visualizations, 21% of the corpus), area charts (2212, 5%), bar charts (3939, 9%), pie charts (1120, 3%), tables (4496, 11%), maps (5182, 13%), dashboards (2472, 6%), and images (7,128, 17%). The remaining 6,248 media (15% of the corpus) did not cluster in thematically coherent ways. Here, we characterize salient elements and trends in these clusters.

Line charts represent the largest cluster of visualizations in our corpus. There are three major substructures: the first comprises line charts depicting the exponential growth of cases in the early stages of the pandemic, and predominantly use log-scales rather than linear scales.

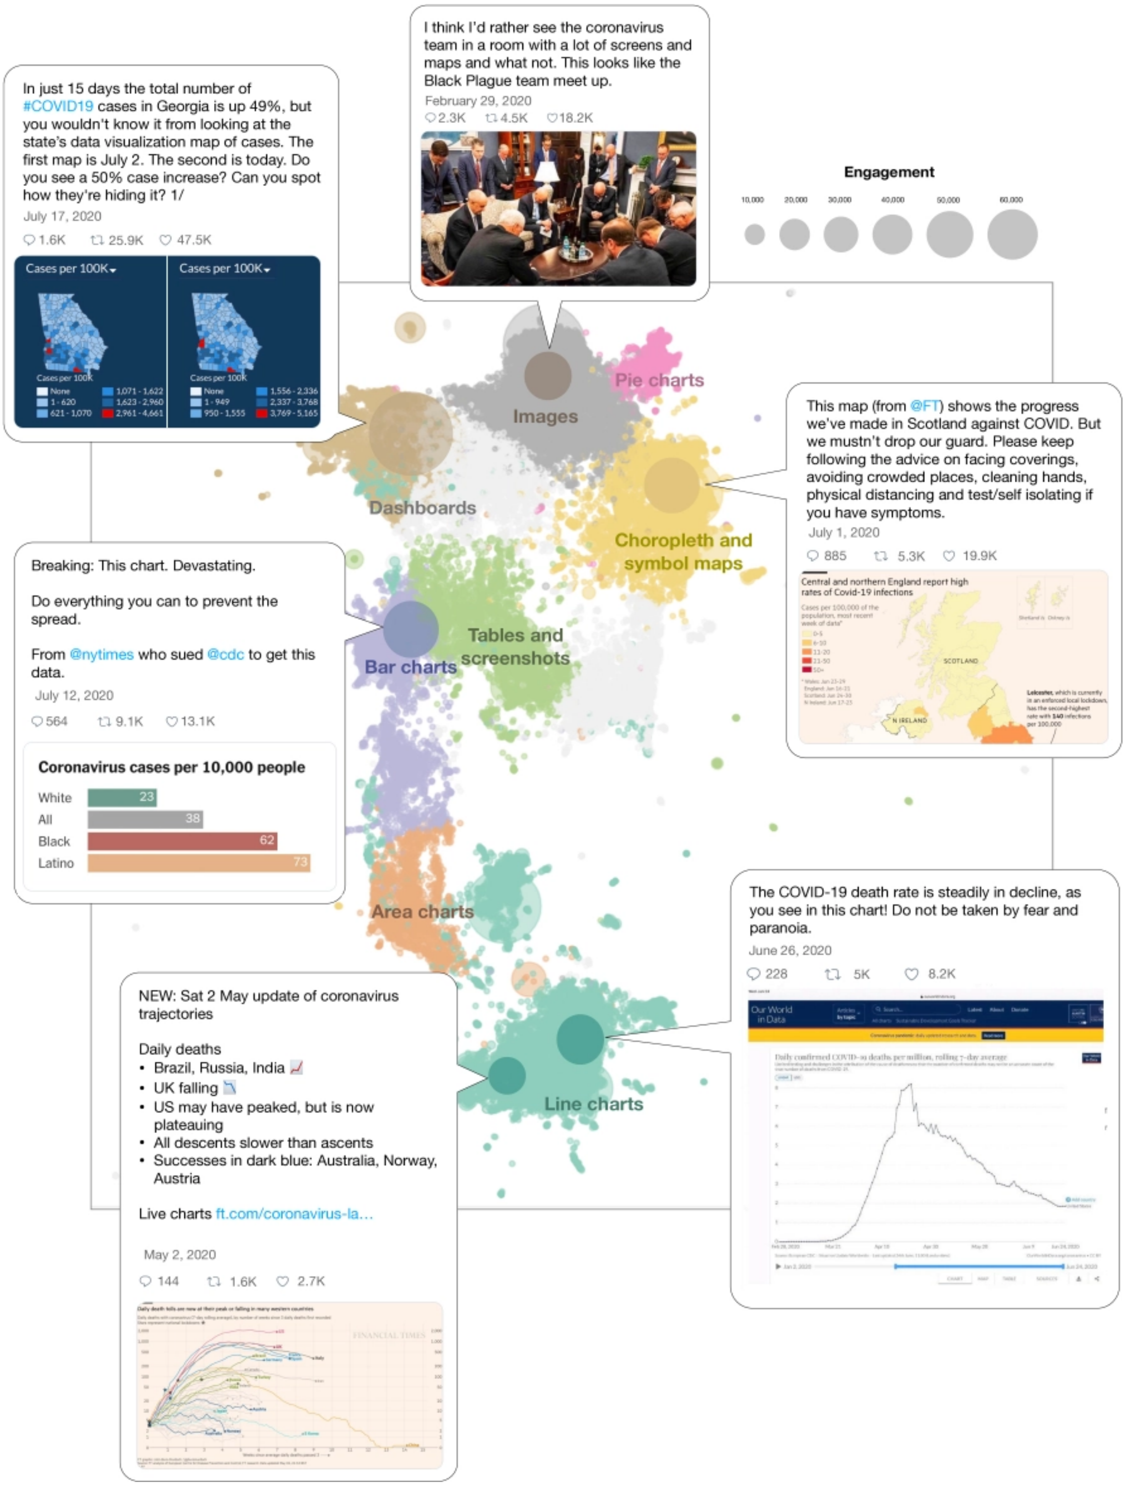

Figure 1. A UMAP visualization of feature embeddings of media found in our Twitter corpus. Color encodes labeled clusters, and size encodes the amount of engagement the media received (i.e., the sum of replies, favorites, retweets, and quote tweets).

Charts from John Burn-Murdoch at the *Financial Times* and charts from the nonprofit Our World in Data are particularly prominent here. A second substructure consists of line charts comparing cases in the United States and the European Union when the US was experiencing its second wave of cases, and the third consists of line charts that visualize economic information. This substructure includes line charts of housing prices, jobs and unemployment, and stock prices (the latter appear to be taken from financial applications and terminals, and often feature additional candlestick marks). Across this cluster, these charts typically depict national or supranational data, include multiple series, and very rarely feature legends or textual annotations (other than labels for each line). Where they do occur, it is to label every point along the lines. Features of the graph are visually highlighted by giving some lines a heavier weight or graying other ones out.

Maps are the second largest cluster of visualizations in our corpus. The overwhelming majority of charts here are choropleths (shaded maps where a geographic region with high COVID rates might be darker, while low-rate regions are lighter). Other visualizations in this cluster include cartograms (the size of a geographic region is proportional its number of COVID infections as a method of comparison) and symbol maps (the size of a circle placed on a geographic region is proportional to COVID infections). The data for these charts span several geographic scales—global trends, country-level data (the US, China, and the UK being particularly salient), and municipal data (states and counties). These maps generally feature heavy annotation including direct labeling of geographic regions with the name and associated data value; arrows and callout boxes also better contextualize the data. For instance, in a widely shared map of the United Kingdom from the *Financial Times*, annotations described how "[t]hree Welsh areas had outbreaks in meatpacking plants in June" and that "Leicester, which is currently in an enforced local lockdown, has the second-highest rate..." These maps depict a wide range of data values including numbers of cases/deaths, metrics normalized per capita, rate of change for cases and/or deaths, mask adherence rates, and the effect of the pandemic on greenhouse gas emissions. Interestingly, choropleth maps of the United States electoral college at both the state- and district-level also appear in the corpus, with the associated tweets comparing the winner of particular regions with the type of pandemic response.

Area charts feature much heavier annotation than line charts (though fewer than maps). Peaks, troughs, and key events (e.g., when lockdowns occurred or when states reopened) are often shaded or labeled with arrows, and line marks are layered to highlight the overall trend or depict the rolling average. When these charts reflect data with a geographic correspondence, this data is often at a more local scale; line charts typically depict national or supranational data, and area charts more often visualized data at the state or county level. Notable subclusters in this group include the viral "Flatten the Curve" graphic, stacked area/streamgraphs, and "skinny bar" charts (charts of temporal data that closely resemble area charts, but use bar marks with narrow widths. Charts from the *New York Times* are especially prominent examples of the latter category—particularly screenshots of a red chart that was featured on the mobile front page.

Bar charts are predominantly encode categorical data and are more consistently and more heavily annotated than area charts. In addition to the annotations described for area charts (direct labeling of the tops of bars, labeled lines and arrows), charts in this cluster often include concise explainer texts.

These texts include some form of extended subtitles, more descriptive axis tick labels, or short passages before or after the bar chart that contextualize the data. Visually, the cluster is equally split between horizontal and vertical charts, and both styles feature a mix of layered, grouped, and stacked bars. Bar chart "races" (e.g., those developed with the Flourish visualization package) are one of the more frequently recurring idioms in this cluster. These are horizontal bar charts depicting the total number of cases per country, and animated over time.

Dashboards and images. While the remaining clusters are thematically coherent, we did not observe as rich a substructure within them. The dashboard cluster is overwhelmingly dominated by screenshots of the Johns Hopkins dashboard, and the image cluster is primarily comprised of reaction memes featuring the photos or caricatures of heads of state.

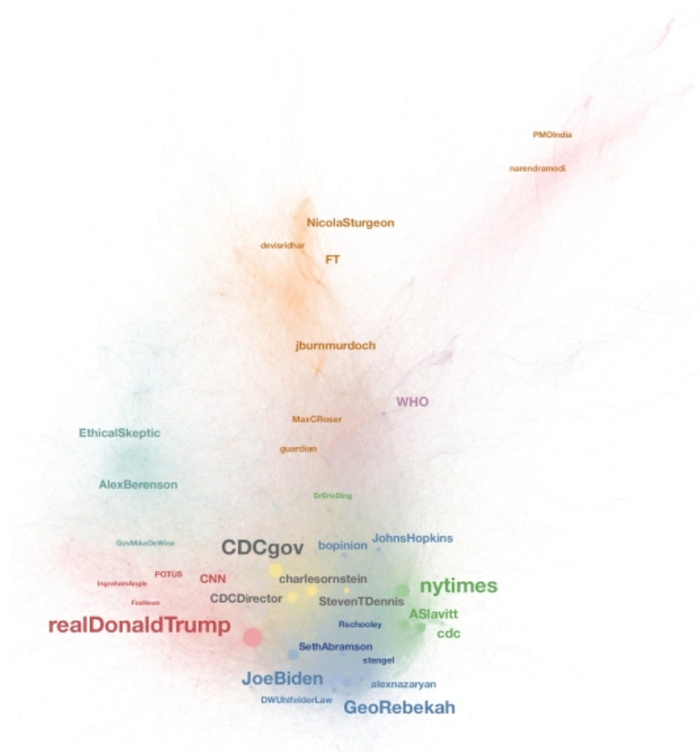

Figure 2: A network visualization of Twitter users appearing in our corpus. Color encodes community as detected by the Louvain method (Blondel et al. 2008), and nodes are sized by their degree of connectedness (i.e., the number of other users they are connected to).

4.1.2 User networks. What are the different networks of Twitter users who share COVID-related data visualizations, and how do they interact with one another? Figure 2 depicts a network graph of Twitter users who discuss (or are discussed) in conversation with the visualizations in Figure 1. This network graph only shows users who are connected to at least five other users (i.e., by replying to them, mentioning them in a tweet, or re-tweeting or quote-tweeting them). The color of each network encodes a specific community as detected by the Louvain method (Blondel et al. 2008), and the graph accounts for the top 10 communities (20,500 users or 72% of the overall graph). We describe the top six networks below listed in order of size (i.e., number of users within each network). While we have designated many of these communities with political orientation (e.g., left- or right-wing), these are only approximations; we recognize that these terms are fundamentally fluid and use them primarily as shorthand to make cross-network comparisons (e.g., mainstream political/media organizations vs. anti-mask protestors).

1. American politics and media (blue). This community features the American center-left, left, mainstream media, and popular or high profile figures (inside and outside of the scientific community). Accounts include politicians (@JoeBiden, @SenWarren), reporters (@joshtpm, @stengel), and public figures like Floridian data scientist @GeoRebekah and actor @GeorgeTakei. The user with the most followers in this community is @neiltyson.

2. American politics and right-wing media (red). This community includes members of the Trump administration, Congress, and right-wing personalities (e.g., @TuckerCarlson). Several accounts of mainstream media organizations also lie within this community (@CNN, @NBCNews), which reflects how often they mention the President (and other government accounts) in their coverage. Notably, these are official organizational accounts rather than those of individual reporters (which mostly show up in the previous group). Several mainstream media organizations are placed equally between these two clusters (@NPR, @washingtonpost). The user with the most followers in this community is @BarackObama.

3. British news media (orange). The largest non-Americentric network roughly corresponds to news media in the UK, with a significant proportion of engagement targeted at the *Financial Times'* successful visualizations by reporter John Burn-Murdoch, as well as coverage of politician Nicola Sturgeon's coronavirus policies. The user with the most followers in this community is @TheEconomist.

4. Anti-mask network (teal). The anti-mask network comprises over 2,500 users (9% of our network graph) and is anchored by former *New York Times* reporter @AlexBerenson, blogger @EthicalSkeptic, and @justin_hart. *The Atlantic*'s @Covid19Tracking project (which collates COVID-19 testing rates and patient outcomes across the United States) and @GovMikeDeWine are also classified as part of this community. Governor DeWine of Ohio is not an anti-masker, but is often the target of anti-mask protest given his public health policies. Anti-mask users also lampoon *The Atlantic*'s project as another example of mainstream misinformation. These dynamics of intertextuality and citation within these networks are especially important here, as anti-mask groups often post screenshots of graphs from "lamestream media" organizations (e.g., *New York Times*) for the purpose of critique and analysis. The user with the most followers in this community is COVID skeptic and billionaire @elonmusk.

5. New York Times-centric network (green). This community is largely an artifact of a single visualization: Andy Slavitt (@ASlavitt), the former acting Administrator of the Centers for Medicare and Medicaid Services under the Obama administration, posted a viral tweet announcing the *New York Times* had sued the CDC (tagged with the incorrect handle @cdc instead of @CDCGov). The attached bar chart showing the racial disparity in COVID cases was shared widely with commentary directly annotated onto the graph itself, or users analyzed the graph through quote-tweets and comments. The user with the most followers in this community is @NYTimes.

6. World Health Organization and health-related news organizations (purple). This community consists of global health organizations, particularly the @WHO and its subsidiary accounts (e.g., @WHOSEARO for Southeast Asia). The user with the most followers in this community is @YouTube.

4.1.3 Descriptive statistics of communities. Table 1 lists summary statistics for the six largest communities in our dataset. There are three statistics of interest: the percentage of verified users (based on the total number of users within a community), the percentage of in-network retweets (based on a community's total number of retweets), and the percentage of original tweets (based on a community's total number of tweets). Twitter verification can often indicate that the account of public interest is authentic (subject to Twitter's black-boxed evaluation standards); it can be a reasonable indication that the account is not a bot. Secondly, a high percentage of in-network retweets can be an indicator of how insular a particular network can be, as it shows how often a user amplifies messages from other in-network members. Finally, the percentage of original tweets shows how much of the content in that particular community is organic (i.e., they write their own content rather than simply amplifying existing work). Communities that have users who use the platform more passively (i.e., they prefer to lurk rather than comment) will have fewer original tweets; communities that have higher levels of active participation will have a higher number of original tweets as a percentage of total tweets.

Table 1. Descriptive Statistics of Communities.

Community #	Verified Users as % of Total Users	In-Network Retweets as % of Total Retweets	Original Tweets as % of Total Tweets
1	8.39	73.30	22.12
2	14.36	75.45	44.75
3	22.92	89.32	34.00
4	10.56	82.17	37.12
5	12.33	58.29	21.57
6	8.94	70.97	37.46

The networks with the highest number of in-network retweets (which can be one proxy for insularity) are the British media (89.32%) and the anti-mask networks (82.17%), and the network with the highest percentage of original tweets is the American politics and right-wing media network (44.75%). Notably, the British news media network has both the *highest* percentage of verified users (22.92%), the *highest* percentage of in-network retweets (89.32%), and the fourth-highest percentage of original tweets (34.00%). As the third largest community in our dataset, we attribute this largely to the popularity of the graphs from the *Financial Times* from a few sources and the constellation of accounts that discussed those visualizations. While other communities (anti-mask, American politics/right-wing media, and WHO/health-related news) shared *more* visualizations, this network shared fewer graphs (1,385) that showed the *second-highest* level of engagement across the six communities (averaging 94 likes, retweets, or mentions per visualization). The network whose visualizations garner the highest level of engagement is the American politics and media network (131 likes, retweets, or mentions per visualization), but they only shared about half of the visualizations (648) compared to their British counterparts.

Through the descriptive statistics, we find that the anti-mask community exhibits very similar patterns to the rest of the networks in our dataset (it has about the same number of users with the same proportion of verified accounts). However, this community has the second highest percentage of in-network retweets (82.17% of all retweets) across the communities, and has the third-highest percentage of original tweets (37.12%, only trailing the World Health Organization network, at 37.46%).

4.1.4 Cross-network comparison of visualization types. Figure 3 depicts the distribution of visualization types by each community, along with descriptive statistics on the numbers of users, charts, and average engagement per tweet. These scatterplots show that there is little variance between the types of visualizations that users in each network share: almost all groups equally use maps or line, area, and bar charts. However, each group usually has one viral visualization—in group 3 (British media), the large yellow circle represents a map from the *Financial Times* describing COVID-19 infections in Scotland; in group 5 (*New York Times*), the large purple circle in the center of the chart represents the viral bar chart from Andy Slavitt describing the racial disparities in COVID cases. The visualizations with the highest number of engagements in each of the six communities is depicted in figure 1.

Overall, we see that each group usually has one viral hit, but that the anti-mask users (group 4) tend to share a wide range of visualizations that garner medium levels of engagement (they have the third-highest number of average engagements in the six communities; an average of 65 likes, shares, and retweets). As a percentage of total tweets, anti-maskers have shared the second-highest number of charts across the top six communities (1,799 charts or 14.75%). They also use the *most* area/line charts and the *least* images across the six communities (images in this dataset usually include memes or photos of politicians). These statistics suggest that anti-maskers tend to be among the most prolific sharers of data visualizations on Twitter, and that they overwhelmingly amplify these visualizations to other users within their network (88.97% of all retweets are in-network).

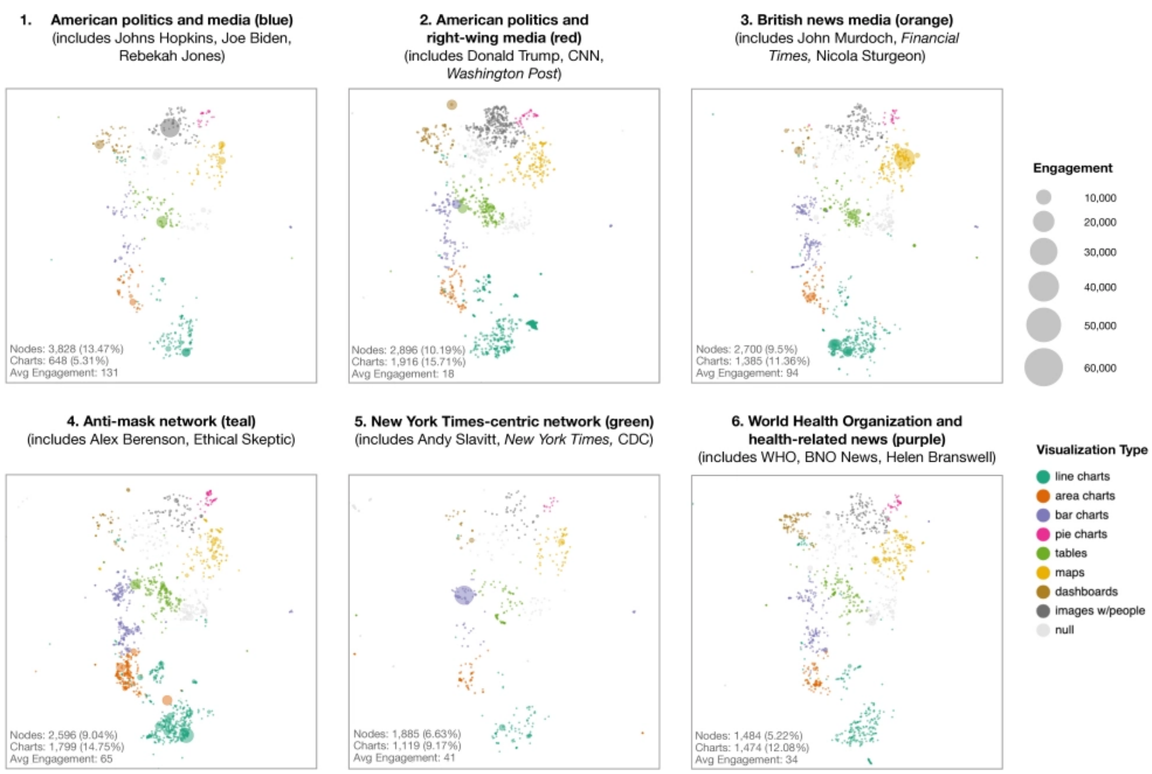

Figure 3: Visualizing the distribution of chart types by network community (with top accounts listed). While every community has produced at least one viral tweet, anti-mask users (group 6) receive higher engagement on average.

4.1.5 Anti-mask visualizations. Figure 4 depicts the data visualizations that are shared by members of the anti-mask network accompanied by a select tweets from each category. While there are certainly visualizations that tend to use a meme-based approach to make their point (e.g., "Hey Fauci...childproof chart!" with the heads of governors used to show the rate of COVID fatalities), many of the visualizations shared by anti-mask Twitter users employ visual forms that are relatively similar to charts that one might encounter at a scientific conference. Many of these tweets use area and line charts to show the discrepancy between the number of projected deaths in previous epidemiological and the numbers of actual fatalities. Others use unit visualizations, tables, and bar charts to compare the severity of coronavirus to the flu. In total, this figure shows the breadth of visualization types that anti-mask users employ to illustrate that the pandemic is exaggerated.

4.2 Anti-mask discourse analysis

The Twitter analysis establishes that anti-maskers are prolific producers and consumers of data visualizations, and that the graphs that they employ are similar to those found in orthodox narratives about the pandemic. Put differently, anti-maskers use "data-driven" narratives to justify their heterodox beliefs. However, a quantitative overview of the visualizations they share and amplify does not in itself help us understand *how* anti-maskers invoke data and scientific reasoning to support policies like re-opening schools and businesses. Anti-maskers are acutely aware that mainstream narratives use data to underscore the pandemic's urgency; they believe that these data sources and visualizations are fundamentally flawed and seek to counteract these biases. This section showcases the different ways that anti-mask groups talk about COVID-related data in discussion forums. What kinds of concerns do they have about the data used to formulate public policies? How do they talk about the limitations of data or create visualizations to convince other members in their physical communities that the pandemic is a hoax?

4.2.1 Emphasis on original content. Many anti-mask users express mistrust for academic and journalistic accounts of the pandemic, proposing to rectify alleged bias by "following the data" and creating their own data visualizations. Indeed, one Facebook group within this study has very strict moderation guidelines that prohibit the sharing of non-original content so that discussions can be "guided solely by the data." Some group administrators even impose news consumption bans on themselves so that "mainstream" models do not "cloud their analysis." In other words, **anti-maskers value unmediated access to information and privilege personal research and direct reading over "expert" interpretations**. While outside content is generally prohibited, Facebook group moderators encourage followers to make their own graphs, which are often shared by prominent members of the group to larger audiences (e.g., on their personal timelines or on other public facing Pages). Particularly in cases where a group or page is led by a few prominent users, follower-generated graphs tend to be highly popular because they often encourage other followers to begin their own data analysis projects, and

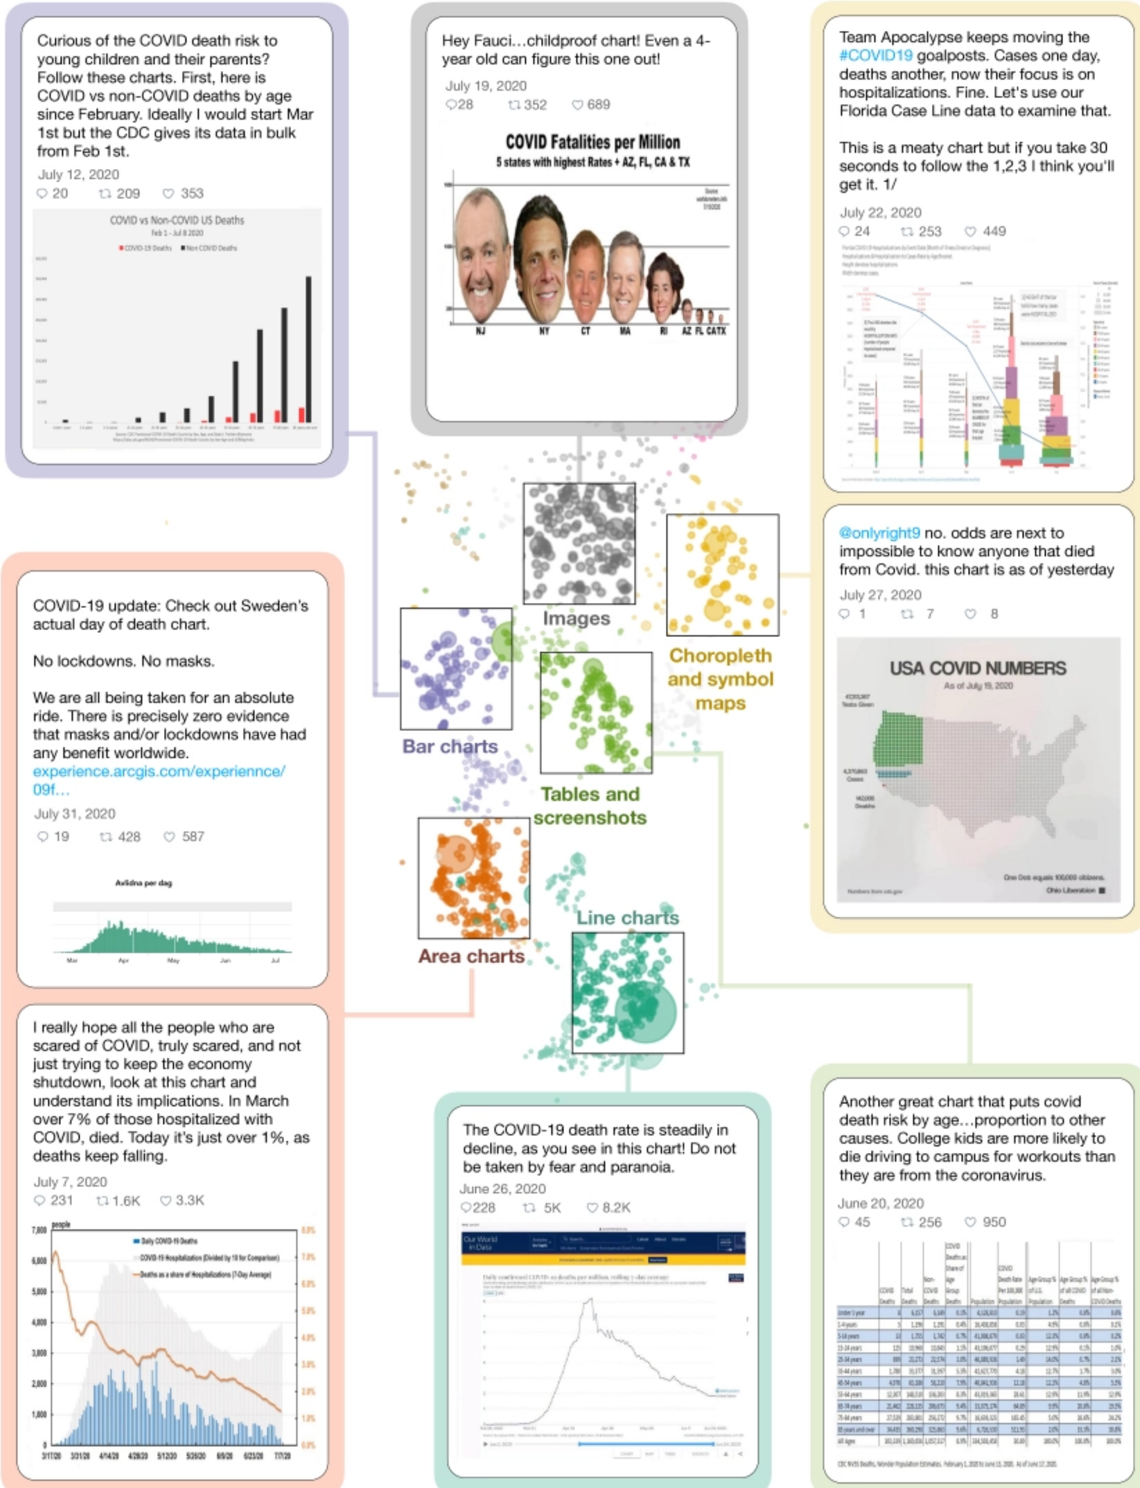

Figure 4: Sample counter-visualizations from the anti-mask user network. While there are meme-based visualizations, anti-maskers on Twitter adopt the same visual vocabulary as visualization experts and the mainstream media.

comments on these posts often deal directly with how to reverse-engineer (or otherwise adjust) the visualization for another locality.

Since some of these groups are focused on a single state (e.g. "Reopen Nevada"), they can fill an information gap: not every county or locality is represented on data dashboards made by local newspapers, health departments, or city governments—if these government entities have dashboards or open data portals at all. In such cases, the emphasis on original content primarily reflects a grassroots effort to ensure access to pandemic-related data where there are no alternatives, and only secondarily serves to constitute an alternative to ideologically charged mainstream narratives. In the rare instances where mainstream visualizations are shared in such a group, it is usually to highlight the ways that mainstream analysis finally matches anti-mask projections, or to show how a journalist, government official, or academic can manipulate the same data source to purposefully mislead readers.

In order to create these original visualizations, users provide numerous tutorials on how to access government health data. These tutorials come either as written posts or as live screencasts, where a user (often a group administrator or moderator) demonstrates the process of downloading information from an open data portal. During the livestream, they editorialize to show which data are the most useful (e.g., *"the data you can download* [from the Georgia Health Department website]*is completely worthless, but the dashboard—which has data that everyday citizens cannot access—actually shows that there are no deaths whatsoever,"* July 13, 2020). Since many local health departments do not have the resources to stand up a new data system specifically for COVID-19, some redirect constituents to the state health department, which may not have granular data for a specific township available for public use. In the absence of data-sharing between states and local governments, users often take it upon themselves to share data with one another (e.g., "[redacted]*brings us this set of data from Minnesota.* [...]*Here it is in raw form, just superimposed on the model,"* May 17, 2020) and they work together to troubleshoot problems with each dataset (e.g., *"thanks. plugging* [sic]*in new .csv file to death dates is frustrating but worth it,"* May 2, 2020).

4.2.2 Critically assessing data sources. Even as these users learn from each other to collect more data, they remain critical about the circumstances under which the data are collected and distributed. Many of the users believe that the most important metrics are missing from government-released data. They express their concerns in four major ways. First, there is an ongoing animated debate within these groups about which metrics matter. Some users contend that deaths, not cases, should be the ultimate arbiter in policy decisions, since case rates are easily "manipulated" (e.g., with increased testing) and do not necessarily signal severe health problems (people can be asymptomatic). The shift in focus is important, as these groups believe that the emphasis on cases and testing often means that rates of COVID deaths by county or township are not reported to the same extent or seriously used for policy making. As one user noted, *"The Alabama public health department doesn't provide deaths per day data* *(that I can tell—you can get it elsewhere). I sent a message asking about that. Crickets so far,"* (July 13, 2020).

Second, users also believe that state and local governments are deliberately withholding data so that they can unilaterally make decisions about whether or not lockdowns are effective. During a Facebook livestream with a Congressional candidate who wanted to "use data for reopening," for example, both the candidate and an anti-mask group administrator discussed the extent to which state executives were willing to obscure the underlying data that were used to justify lockdown procedures (August 30, 2020). To illustrate this, the candidate emphasized a press conference in which journalists asked the state executive whether they would consider making the entire contact tracing process public, which would include releasing the name of the bar where the outbreak started. In response, the governor argued that while transparency about the numbers were important, the state would not release the name of the bar, citing the possibility of stigmatization and an erosion of privacy. This soundbite—*"we have the data, but we won't give it to you"*—later became a rallying cry for anti-mask groups in this state following this livestream. *"I hate that they're not being transparent in their numbers and information they're giving out,"* another user wrote. *"They need to be honest and admit they messed up if it isn't as bad as they're making it out to be.* [...]*We need honesty and transparency."*

This plays into a third problem that users identify with the existing data: that datasets are constructed in fundamentally subjective ways. They are coded, cleaned, and aggregated either by government data analysts with nefarious intentions or by organizations who may not have the resources to provide extensive documentation. *"Researchers can define their data set anyhow* [sic]*they like in absence of generally accepted (preferably specified) definitions,"* one user wrote on June 23, 2020. *"Coding data is a big deal—and those definitions should be offered transparently by every state. Without a national guideline—we are left with this mess."* The lack of transparency within these data collection systems—which many of these users infer as a lack of honesty—erodes these users' trust within both government institutions and the datasets they release.

Even when local governments do provide data, however, these users also contend that the data requires context in order for any interpretation to be meaningful. For example, metrics like hospitalization and infection rates are still *"vulnerable to all sorts of issues that make [these] data less reliable than deaths data"* (June 23, 2020), and require additional inquiry before the user considers making a visualization. In fact, there are multiple threads every week where users debate how representative the data are of the population given the increased rate of testing across many states. For some users, random sampling is the only way to really know the true infection rate, as (1) testing only those who show symptoms gives us an artificially high infection rate, and (2) testing asymptomatic people tells us what we already know—that the virus is not a threat. These groups argue that the conflation of asymptomatic and symptomatic cases therefore makes it difficult for anyone to actually determine the severity of the pandemic. *"We are counting 'cases' in ways we never did for any other virus,"* a user writes, *"and we changed how we counted in the middle of the game. It's classic garbage in, garbage out at this point. If it could be clawed back to ONLY symptomatic and/or contacts, it could be a useful guide* [for comparison], *but I don't see that happening"* (August 1, 2020).

Similarly, these groups often question the context behind measures like "excess deaths." While the CDC has provided visualizations that estimate the number of excess deaths by week (CDC 2020), users take screenshots of the websites and debate whether or not they can be attributed to the coronavirus. *"You can't simply subtract the current death tally from the typical value for this time of year and attribute the difference to Covid,"* a user wrote. *"Because of the actions of our governments, we are actually causing excess deaths. Want to kill an old person quickly? Take away their human interaction and contact. Or force them into a rest home with other infected people. Want people to die from preventable diseases? Scare them away from the hospitals, and encourage them to postpone their medical screenings, checkups, and treatments* [...]*The numbers are clear. By trying to mitigate one problem, we are creating too many others, at too high a price"* (September 5, 2020).

4.2.3 Critically assessing data representations. Even beyond downloading datasets from local health departments, users in these groups are especially attuned to the ways that specific types of visualizations can obscure or highlight information. In response to a visualization where the original poster (OP) created a bar chart of death counts by county, a user commented: *"the way data is presented can also show bias. For example in the state charts, counties with hugely different populations can be next to each other. The smaller counties are always going to look calm even if per capita they are doing the same or worse. Perhaps you could do a version of the charts where the hardest hit county is normalized per capita to 1 and compare counties that way,"* to which the OP responded, *"it is never biased to show data in its entirety, full scale"* (August 14, 2020).

An ongoing topic of discussion is whether to visualize absolute death counts as opposed to deaths per capita, and it is illustrative of a broader mistrust of mediation. For some, "raw data" (e.g., counts) provides more accurate information than any data transformation (e.g., death rate per capita, or even visualizations themselves). For others, screenshots of tables are the most faithful way to represent the data, so that people can see and interpret it for themselves. *"No official graphs,"* said one user. *"Raw data only. Why give them an opportunity to spin?"* (June 14, 2020). These users want to understand and analyze the information for themselves, free from biased, external intervention.

4.2.4 Identifying bias and politics in data. While users contend that their data visualizations objectively illustrate how the pandemic is no worse than the flu, they are similarly mindful to note that these analyses only represent partial perspectives that are subject to individual context and interpretation. *“I've never claimed to have no bias. Of course I am biased, I'm human,”* says one prolific producer of anti-mask data visualizations. *“That's why scientists use controls... to protect ourselves from our own biases. And this is one of the reasons why I disclose my biases to you. That way you can evaluate my conclusions in context. Hopefully, by staying close to the data, we keep the effect of bias to a minimum”* (August 14, 2020). They are ultimately mindful of the subjectivity of human interpretation, which leads them to analyzing the data for themselves.

More tangibly, however, these groups seek to identify bias by being critical about specific profit motives that come from releasing (or suppressing) specific kinds of information. Many of the users within these groups are skeptical about the potential benefits of a coronavirus vaccine, and as a point of comparison, they often reference how the tobacco industry has historically manipulated science to mislead consumers. These groups believe that pharmaceutical companies have similarly villainous profit motives, which leads the industry to inflate data about the pandemic in order to stoke demand for a vaccine. As one user lamented, *“I wish more of the public would do some research into them and see how much of a risk they are but sadly most wont* [sic]*—because once you do and you see the truth on them, you get labeled as an ‘antivaxxer’ which equates to fool. In the next few years, the vaccine industry is set to be a nearly 105 billion dollar industry. People should really consider who profits off of our ignorance”* (August 24, 2020).

4.2.5 Appeals to scientific authority. Paradoxically, these groups also seek ways to validate their findings through the scientific establishment. Many users prominently display their scientific credentials (e.g., referring to their doctoral degrees or prominent publications in venues like Nature) which uniquely qualify them as insiders who are most well-equipped to criticize the scientific community. Members who perform this kind of expertise often point to 2013 Nobel Laureate Michael Levitt's assertion that lockdowns do nothing to save lives [67] as another indicator of scientific legitimacy. Both Levitt and these anti-mask groups identify the dangerous convergence of science and politics as one of the main barriers to a more reasonable and successful pandemic response, and they construct their own data visualizations as a way to combat what they see as health misinformation. *“To be clear. I am not downplaying the COVID epidemic,”* said one user. *“I have never denied it was real. Instead, I've been modeling it since it began in Wuhan, then in Europe, etc.* [...]*What I have done is follow the data. I've learned that governments, that work for us, are too often deliberately less than transparent when it comes to reporting about the epidemic”* (July 17, 2020). For these anti-mask users, their approach to the pandemic is grounded in a *more* scientific rigor, not less.

4.2.6 Developing expertise and processes of critical engagement. The goal of many of these groups is ultimately to develop a network of well-informed citizens engaged in analyzing data in order to make measured decisions during a global pandemic. *“The other side says that they use evidence-based medicine to make decisions,”* one user wrote, *“but the data and the* *science do not support current actions"* (August 30, 2020). The discussion-based nature of these Facebook groups also give these followers a space to learn and adapt from others, and to develop processes of critical engagement. Long-time followers of the group often give small tutorials to new users on how to read and interpret specific visualizations, and users often give each other constructive feedback on how to adjust their graphic to make it more legible or intuitive. Some questions and comments would not be out of place at all at a visualization research poster session: *"This doesn't make sense. What do the colors mean? How does this demonstrate any useful information?"* (July 21, 2020) These communities use data analysis as a way to socialize and enculturate their users; they promulgate data literacy practices as a way of inculcating heterodox ideology. The transmission of data literacy, then, becomes a method of political radicalization.

These individuals as a whole are extremely willing to help others who have trouble interpreting graphs with multiple forms of clarification: by helping people find the original sources so that they can replicate the analysis themselves, by referencing other reputable studies that come to the same conclusions, by reminding others to remain vigilant about the limitations of the data, and by answering questions about the implications of a specific graph. The last point is especially salient, as it surfaces both what these groups see as a reliable measure of how the pandemic is unfolding and what they believe they should do with the data. These online communities therefore act as a sounding board for thinking about how best to effectively mobilize the data towards more measured policies like slowly reopening schools. *"You can tell which places are actually having flare-ups and which ones aren't,"* one user writes. *"Data makes us calm."* (July 21, 2020)

Additionally, **followers in these groups also use data analysis as a way of bolstering social unity and creating a community of practice.** While these groups highly value scientific expertise, they also see collective analysis of data as a way to bring communities together within a time of crisis, and being able to transparently and dispassionately analyze the data is crucial for democratic governance. In fact, the explicit motivation for many of these followers is to find information so that they can make the best decisions for their families—and by extension, for the communities around them. *"Regardless of your political party, it is incumbent on all of us to ask our elected officials for the data they use to make decisions,"* one user said during a live streamed discussion. *"I'm speaking to you as a neighbor: request the data.* [...]*As a Mama Bear, I don't care if Trump says that it's okay, I want to make a decision that protects my kids the most. This data is especially important for the moms and dads who are concerned about their babies"* (August 30, 2020). As Kate Starbird et al. (2019) have demonstrated, strategic information operations require the participation of online communities to consolidate and amplify these messages: these messages become powerful when emergent, organic crowds (rather than hired trolls and bots) iteratively contribute to a larger community with shared values and epistemologies.

Group members repost these analyses onto their personal timelines to start conversations with friends and family in hopes that they might be able to congregate in person. However, many of these conversations result in frustration. *"I posted virus data from the CDC, got into discussion with people and in the end several straight out voiced they had no interest in the data,"* one user sighed. *"My post said 'Just the facts.'* [screenshot from the CDC]*People are emotionally invested in their beliefs and won't be swayed by data. It's disturbing"* (August 14, 2020). Especially when these conversations go poorly, followers solicit advice from each other about how to move forward when their children's schools close or when family members do not "follow the data." One group even organized an unmasked get-together at a local restaurant where they passed out t-shirts promoting their Facebook group, took selfies, and discussed a lawsuit that sought to remove their state's emergency health order (September 12, 2020). The lunch was organized such that the members who wanted to first attend a Trump rolling rally could do so and *"drop in afterward for some yummy food and fellowship"* (September 8, 2020).

4.2.7 Applying data to real-world situations. Ultimately, anti-mask users emphasize that they need to apply this data to real-world situations. The same group that organized the get-together also regularly hosts live-streams with guest speakers like local politicians, congressional candidates, and community organizers, all of whom instruct users on how to best agitate for change armed with the data visualizations shared in the group. *"You're a mom up the street, but you're not powerless,"* emphasized one of the guest speakers. *"Numbers matter! What is just and what is true matters.* [...]*Go up and down the ladder—start real local. Start with the lesser magistrates, who are more accessible, easier to reach, who will make time for you."* (July 23, 2020)

These groups have been incredibly effective at galvanizing a network of engaged citizens towards concrete political action. Local officials have relied on data narratives generated in these groups to call for a lawsuit against the Ohio Department of Health (July 20, 2020). In Texas, a coalition of mayors, school board members, and city council people investigated the state's COVID-19 statistics and discovered that a backlog of unaudited tests was distorting the data, prompting Texas officials to employ a forensic data team to investigate the surge in positive test rates (Carroll 2020). *"There were over a million pending assignments* [that were distorting the state's infection rate]," the city councilperson said to the group's 40,000+ followers. *"We just want to make sure that the information that is getting out there is giving us the full picture."* (August 17, 2020) Another Facebook group solicited suggestions from its followers on how to support other political groups who need data to support lawsuits against governors and state health departments. *"If you were suddenly given access to all the government records and could interrogate any official,"* a group administrator asked, *"what piece of data or documentation would you like to inspect?"* (September 11, 2020) The message that runs through these threads is unequivocal: that data is the only way to set fear-bound politicians straight, and using better data is a surefire way towards creating a safer community.

5 Discussion

Anti-maskers have deftly used social media to constitute a cultural and discursive arena devoted to addressing the pandemic and its fallout through practices of data literacy. Data literacy is a quintessential criterion for membership within the community they have created. The prestige of both individual anti-maskers and the larger Facebook groups to which they belong is tied to displays of skill in accessing, interpreting, critiquing, and visualizing data, as well as the pro-social willingness to share those skills with other interested parties. This is a community of practice (Lave and Wenger 1991, Wenger 1998) focused on acquiring and transmitting expertise, and on translating that expertise into concrete political action. Moreover, this is a subculture shaped by mistrust of established authorities and orthodox scientific viewpoints. Its members value individual initiative and ingenuity, trusting scientific analysis only insofar as they can replicate it themselves by accessing and manipulating the data firsthand. They are highly reflexive about the inherently biased nature of any analysis, and resent what they view as the arrogant self-righteousness of scientific elites.

As a subculture, anti-masking amplifies anti-establishment currents pervasive in U.S. political culture. Data literacy, for anti-maskers, exemplifies distinctly American ideals of intellectual self-reliance, which historically takes the form of rejecting experts and other elites (Hofstadter 1966). The counter-visualizations that they produce and circulate not only challenge scientific consensus, but they also assert the value of independence in a society that they believe promotes an overall de-skilling and dumbing-down of the population for the sake of more effective social control (Elisha 2011, Hochschild 2016, Tripodi 2018). As they see it, to counter-visualize is to engage in an act of resistance against the stifling influence of central government, big business, and liberal academia. Moreover, their simultaneous appropriation of scientific rhetoric and rejection of scientific authority also reflects longstanding strategies of Christian fundamentalists seeking to challenge the secularist threat of evolutionary biology (Bielo 2019).

So how do these groups diverge from scientific orthodoxy if they are using the same data? We have identified a few sleights of hand that contribute to the broader epistemological crisis we identify between these groups and the majority of scientific researchers. For instance, they argue that there is an outsized emphasis on *deaths* versus cases: if the current datasets are fundamentally subjective and prone to manipulation (e.g., increased levels of faulty testing, asymptomatic vs. symptomatic cases), then deaths are the only reliable markers of the pandemic's severity. Even then, these groups believe that deaths are an additionally problematic category because doctors are using a COVID diagnosis as the main cause of death (i.e., people who die because of COVID) when in reality there are other factors at play (i.e., dying with but not because of COVID). Since these categories are fundamentally subject to human interpretation, especially by those who have a vested interest in reporting as many

COVID deaths as possible, these numbers are vastly over-reported, unreliable, and no more significant than the flu.

Another point of contention is that of lived experience: in many of these cases, users do not themselves know a person who has experienced COVID, and the statistics they see on the news show the severity of the pandemic in vastly different parts of the country. Since they do not see their experience reflected in the narratives they consume, they look for hyperlocal data to help guide their decision-making. But since many of these datasets do not always exist on such a granular level, this information gap feeds into a larger social narrative about the government's suppression of critical data and the media's unwillingness to substantively engage with the subjectivity of coronavirus data reporting.

Most fundamentally, the groups we studied believe that **science is a process, and not an institution**. As we have outlined in the case study, these groups mistrust the scientific establishment ("Science") because they believe that the institution has been corrupted by profit motives and politics. The knowledge that the CDC and academics have created cannot be trusted because they need to be subject to increased doubt, and not accepted as consensus. In the same way that climate change skeptics have appealed to Karl Popper's theory of falsification to show why climate science needs to be subjected to continuous scrutiny in order to be valid (Fischer 2019), we have found that anti-mask groups point to Thomas Kuhn's *The Structure of Scientific Revolutions* to show how their anomalous evidence—once dismissed by the scientific establishment—will pave the way to a new paradigm (*"As I've recently described, I'm no stranger to presenting data that are inconsistent with the narrative. It can get ugly. People do not give up their paradigms easily.* [...]*Thomas Kuhn wrote about this phenomenon, which occurs repeatedly throughout history. Now is the time to hunker down. Stand with the data,"* August 5, 2020). For anti-maskers, valid science must be a process they can critically engage for themselves in an unmediated way. *Increased* doubt, not consensus, is the marker of scientific certitude.

Arguing that anti-maskers simply need more scientific literacy is to characterize their approach as uninformed and inexplicably extreme. This study shows the opposite: users in these communities are deeply invested in forms of critique and knowledge production that they recognize as markers of scientific expertise. If anything, anti-mask science has extended the traditional tools of data analysis by taking up the theoretical mantle of recent critical studies of visualization (Correll 2019; D'Ignzaio and Klein 2020). Anti-mask approaches acknowledge the subjectivity of how datasets are constructed, attempt to reconcile the data with lived experience, and these groups seek to make the process of understanding data as transparent as possible in order to challenge the powers that be. For example, one of the most popular visualizations within the Facebook groups we studied were unit visualizations, which are popular among anti-maskers and computer scientists for the same reasons: they provide more information, better match a reader's mental model, and they allow users to interact with them in new and more interesting ways (Park et al. 2018). Barring tables, they are the most unmediated way to interact with data: one dot represents one person.

Similarly, these groups' impulse to mitigate bias and increase transparency (often by dropping the use of data they see as "biased") echoes the organizing ethos of computer science research that seeks to develop "technological solutions regarding potential bias" or "ground research on fairness, accountability, and transparency" (Association for Computing Machinery 2020). In other words, these groups see themselves as engaging deeply within multiple aspects of the scientific process—interrogating the datasets, analysis, and conclusions—and still university researchers might dismiss them in leading journals as "scientifically illiterate" (Miller 2020). In an interview with the Department of Health and Human Services podcast, even Anthony Fauci (Chief Medical Advisor to the US President) noted: *"one of the problems we face in the United States is that unfortunately, there is a combination of an anti-science bias* [...]*people are, for reasons that sometimes are, you know, inconceivable and not understandable, they just don't believe science"* (Fauci and Caputo 2020).

We use Dr. Fauci's provocation to illustrate how understanding the way that anti-mask groups think about science is crucial to grappling with the contested state of expertise in American democracy. In a study of Tea Party supporters in Louisiana, Arlie Russell Hochschild (Hochschild 2016) explains the intractable partisan rift in American politics by emphasizing the importance of a "deep story": a subjective prism that people use in order to make sense of the world and guide the way they vote. For Tea Party activists, this deep story revolved around anger towards a federal system ruled by liberal elites who pander to the interests of ethnic and religious minorities, while curtailing the advantages that White, Christian traditionalists view as their American birthright. We argue that the anti-maskers' deep story draws from similar wells of resentment, but adds a particular emphasis on the usurpation of scientific knowledge by a paternalistic, condescending elite that expects intellectual subservience rather than critical thinking from the lay public.

To be clear, we are not promoting these views. Instead, we seek to better understand how data literacy, as a both a set of skills and a moral virtue championed within academic computer science, can take on distinct valences in different cultural contexts. A more nuanced view of data literacy, one that recognizes multiplicity rather than uniformity, offers a more robust account of how data visualization circulates in the world. This culturally and socially situated analysis demonstrates why increasing access to raw data or improving the informational quality of data visualizations is not sufficient to bolster public consensus about scientific findings. Projects that examine the cognitive basis of visualization or seek to make "better" or "more intuitive" visualizations (Kosara 2016) will not meaningfully change this phenomenon: anti-mask protestors already use visualizations, and do so extremely effectively. Moreover, in emphasizing the politicization of pandemic data, our account helps to explain the striking correlation between practices of counter-visualization and the politics of anti-masking.

For members of this social movement, counter-visualization and anti-masking are complementary aspects of resisting the tyranny of institutions that threaten to usurp individual liberties to think freely and act accordingly.

6 Implications and Conclusion

This paper has investigated anti-mask counter-visualizations on social media in two ways: quantitatively, we identify the main types of visualizations that are present within different networks (e.g., pro- and anti-mask users), and we show that anti-mask users are prolific and skilled purveyors of data visualizations. These visualizations are popular, use orthodox visualization methods, and are promulgated as a way to convince others that public health measures are unnecessary. In our qualitative analysis, we use an ethnographic approach to illustrate how COVID counter-visualizations actually reflect a deeper epistemological rift about the role of data in public life, and that the practice of making counter-visualizations reflects a participatory, heterodox approach to information sharing. Convincing anti-maskers to support public health measures in the age of COVID-19 will require more than “better” visualizations, data literacy campaigns, or increased public access to data. Rather, it requires a sustained engagement with the social world of visualizations and the people who make or interpret them.

While academic science is traditionally a system for producing knowledge within a laboratory, validating it through peer review, and sharing results within subsidiary communities, anti-maskers reject this hierarchical social model. They espouse a vision of science that is radically egalitarian and individualist. This study forces us to see that coronavirus skeptics champion science as a personal practice that prizes rationality and autonomy; for them, it is *not* a body of knowledge certified by an institution of experts. Calls for data or scientific literacy therefore risk recapitulating narratives that anti-mask views are the product of individual ignorance rather than coordinated information campaigns that rely heavily on networked participation. Recognizing the *systemic* dynamics that contribute to this epistemological rift is the first step towards grappling with this phenomenon, and the findings presented in this paper corroborate similar studies about the impact of fake news on American evangelical voters (Tripodi 2018) and about the limitations of fact-checking climate change denialism (Fischer 2019).

Calls for media literacy—especially as an ethics smokescreen to avoid talking about larger structural problems like white supremacy—are problematic when these approaches are deficit-focused and trained primarily on individual responsibility. Powerful research and media organizations paid for by the tobacco or fossil fuel industries (Oreskes and Conway 2010; Proctor 2011) have historically capitalized on the skeptical impulse that the “science simply isn't settled,” prompting people to simply “think for themselves” to horrifying ends. The attempted coup on January 6, 2021 has similarly illustrated that well-calibrated, well-funded systems of coordinated disinformation can be particularly dangerous when they are *designed* to appeal to skeptical people. While individual insurrectionists are no doubt to blame for their own acts of violence, the coup relied on a collective effort fanned by people questioning, interacting, and sharing these ideas with other people. These skeptical narratives are powerful because they resonate with these people's lived experience and—crucially—because they are posted by influential accounts across influential platforms.

Broadly, the findings presented in this paper also challenge conventional assumptions in human-computer interaction research about who imagined users might be: visualization experts traditionally design systems for scientists, business analysts, or journalists. Researchers create systems intended to democratize processes of data analysis and inform a broader public about how to use data, often in the clean, sand-boxed environment of an academic lab. However, this literature often focuses narrowly on promoting expressivity (either of current or new visualization techniques), assuming that improving visualization tools will lead to improving public understanding of data. This paper presents a community of users that researchers might not consider in the systems building process (i.e., supposedly "data illiterate" anti-maskers), and we show how the binary opposition of literacy/illiteracy is insufficient for describing how orthodox visualizations can be used to promote unorthodox science. Understanding how these groups skillfully manipulate data to undermine mainstream science requires us to adjust the theoretical assumptions in HCI research about how data can be leveraged in public discourse.

What, then, are visualization researchers and social scientists to do? One step might be to grapple with the social and political dimensions of visualizations at the *beginning*, rather than the end, of projects (Correll 2019. This involves in part a shift from positivist to interpretivist frameworks in visualization research, where we recognize that knowledge we produce in visualization systems is fundamentally "multiple, subjective, and socially constructed" (Meyer and Dykes 2019). A secondary issue is one of uncertainty: Jessica Hullman and Zeynep Tufekci (among others) have both showed how *not* communicating the uncertainty inherent in scientific writing has contributed to the erosion of public trust in science (Hullmann et al. 2019, Tufekci 2020). As Tufekci demonstrates (and our data corroborates), the CDC's initial public messaging that masks were ineffective—followed by a quick public reversal—seriously hindered the organization's ability to effectively communicate as the pandemic progressed. As we have seen, people are not simply passive consumers of media: anti-mask users *in particular* were predisposed to digging through the scientific literature and highlighting the uncertainty in academic publications that media organizations elide. When these uncertainties did not surface within public-facing versions of these studies, people began to assume that there was a broader cover-up (Tufekci 2020).

But as Hullman shows, there are at least two major reasons why uncertainty *hasn't* traditionally been communicated to the public (Hullman 2020). Researchers often do not believe that people will understand and be able to interpret results that communicate uncertainty (which, as we have shown, is a problematic assumption at best). However, visualization researchers also do not have a robust body of understanding about *how*, and when, to communicate uncertainty (let alone how to do so effectively). There are exciting threads of visualization research that investigate how users' interpretive frameworks can change the overarching narratives they glean from the data (Hullman and Diakopoulos 2011; Peck et al. 2019; Segel and Heer 2010). Instead of championing absolute certitude or objectivity, this research pushes us to ask how scientists and visualization researchers alike might express uncertainty in the data so as to recognize its socially and historically situated nature.

In other words, our paper introduces new ways of thinking about "democratizing" data analysis and visualization. Instead of treating increased adoption of data-driven storytelling as an unqualified good, we show that data visualizations are not simply tools that people use to understand the epidemiological events around them. They are a battleground that highlight the contested role of expertise in modern American life.

References for Chapter 2

Adar, Eytan, and Elsie Lee. "Communicative Visualizations as a Learning Problem." *IEEE Transactions on Visualization and Computer Graphics*, 2020, 1–11. https://doi.org/10.1109/TVCG.2020.3030375.

Ahearn, Laura M. "Literacy, Power, and Agency: Love Letters and Development in Nepal." *Language and Education* 18, no. 4 (September 2, 2004): 305–16. https://doi.org/10.1080/09500780408666883.

Alberda, Alexandra P. "COVID-19 Data Literacy Is for Everyone." Medium, May 5, 2020. https://medium.com/nightingale/covid-19-data-literacy-is-for-everyone-46120b58cec9.

Alper, Basak, Nathalie Henry Riche, Fanny Chevalier, Jeremy Boy, and Metin Sezgin. "Visualization Literacy at Elementary School." In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 5485–97. Denver Colorado USA: ACM, 2017. https://doi.org/10.1145/3025453.3025877.

Anti-Eviction Mapping Project. *(Dis)Location Black Exodus*, 2019. http://archive.org/details/dislocationblackexodus.

Arif, Ahmer, Leo Graiden Stewart, and Kate Starbird. "Acting the Part: Examining Information Operations Within #BlackLivesMatter Discourse." *Proceedings of the ACM on Human-Computer Interaction* 2, no. CSCW (November 2018): 1–27. https://doi.org/10.1145/3274289.

Association for Computing Machinery. "ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT)," September 8, 2020. https://facctconference.org/.

Bail, Christopher A., Lisa P. Argyle, Taylor W. Brown, John P. Bumpus, Haohan Chen, M. B. Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. "Exposure to Opposing Views on Social Media Can Increase Political Polarization." *Proceedings of the National Academy of Sciences* 115, no. 37 (September 11, 2018): 9216–21. https://doi.org/10.1073/pnas.1804840115.

Baym, Nancy K. "Data Not Seen: The Uses and Shortcomings of Social Media Metrics." *First Monday* 18, no. 10 (October 7, 2013). https://doi.org/10.5210/fm.v18i10.4873.

Berriche, Manon, and Sacha Altay. "Internet Users Engage More with Phatic Posts than with Health Misinformation on Facebook." *Palgrave Communications* 6, no. 1 (April 28, 2020): 1–9. https://doi.org/10.1057/s41599-020-0452-1.

Bielo, James S. "'Particles-to-People…Molecules-to-Man': Creationist Poetics in Public Debates." *Journal of Linguistic Anthropology* 29, no. 1 (2019): 4–26. https://doi.org/10.1111/jola.12205.

Blondel, Vincent D., Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. "Fast Unfolding of Communities in Large Networks." *Journal of Statistical Mechanics: Theory and Experiment* 2008, no. 10 (October 2008): P10008. https://doi.org/10.1088/1742-5468/2008/10/P10008.

Börner, Katy, Andreas Bueckle, and Michael Ginda. "Data Visualization Literacy: Definitions, Conceptual Frameworks, Exercises, and Assessments." *Proceedings of the National Academy of Sciences* 116, no. 6 (February 5, 2019): 1857–64. https://doi.org/10.1073/pnas.1807180116.

Börner, Katy, Adam Maltese, Russell Nelson Balliet, and Joe Heimlich. "Investigating Aspects of Data Visualization Literacy Using 20 Information Visualizations and 273 Science Museum Visitors." *Information Visualization* 15, no. 3 (July 1, 2016): 198–213. https://doi.org/10.1177/1473871615594652.

Bowe, Emily, Erin Simmons, and Shannon Mattern. "Learning from Lines: Critical COVID Data Visualizations and the Quarantine Quotidian." *Big Data & Society*, July 27, 2020. https://doi.org/10.1177/2053951720939236.

Bowker, Geoffrey C, and Susan Leigh Star. *Sorting Things Out: Classification and Its Consequences*. Cambridge, MA: MIT Press, 2000.

Boy, Jeremy, Ronald A. Rensink, Enrico Bertini, and Jean-Daniel Fekete. "A Principled Way of Assessing Visualization Literacy." *IEEE Transactions on Visualization and Computer Graphics* 20, no. 12 (December 31, 2014): 1963–72. https://doi.org/10.1109/TVCG.2014.2346984.

boyd, danah. "You Think You Want Media Literacy… Do You?" Medium, March 16, 2018. https://points.datasociety.net/you-think-you-want-media-literacy-do-you-7cad6af18ec2.

Buckingham, David. "Fake News: Is Media Literacy the Answer?" *David Buckingham* (blog), January 12, 2017. https://davidbuckingham.net/2017/01/12/fake-news-is-media-literacy-the-answer/.

Buolamwini, Joy, and Timnit Gebru. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." *Proceedings of Machine Learning Research* 81 (2018): 1–15.

Burgess, Jean, and Ariadna Matamoros-Fernández. "Mapping Sociocultural Controversies across Digital Media Platforms: One Week of #gamergate on Twitter, YouTube, and

Tumblr.” *Communication Research and Practice* 2, no. 1 (January 2, 2016): 79–96. https://doi.org/10.1080/22041451.2016.1155338.

Calero Valdez, André, Martina Ziefle, and Michael Sedlmair. “Studying Biases in Visualization Research: Framework and Methods.” In *Cognitive Biases in Visualizations*, edited by Geoffrey Ellis, 13–27. Cham: Springer International Publishing, 2018. https://doi.org/10.1007/978-3-319-95831-6_2.

Campolo, Alexander. “Flattening the Curve; Visualization and Pandemic Knowledge.” *Formations* (blog), April 1, 2020. https://sifk.uchicago.edu/news/flattening-the-curve-visualization-and-pandemic-knowledge/.

Carroll, Joe. “Texas Governor Says Positive-Rate Surge Under Investigation.” *Bloomberg*, August 13, 2020. https://www.bloomberg.com/news/articles/2020-08-13/texas-governor-says-positivity-rate-surge-under-investigation.

CDC. “Coronavirus Disease 2019 (COVID-19).” Centers for Disease Control and Prevention, February 11, 2020. https://www.cdc.gov/coronavirus/2019-ncov/prevent-getting-sick/cloth-face-cover-guidance.html.

Charmaz, Kathy. *Constructing Grounded Theory*. London ; Thousand Oaks, Calif: Sage Publications, 2006.

Chen, Emily, Kristina Lerman, and Emilio Ferrara. “Tracking Social Media Discourse About the COVID-19 Pandemic: Development of a Public Coronavirus Twitter Data Set.” *JMIR Public Health and Surveillance* 6, no. 2 (2020): e19273. https://doi.org/10.2196/19273.

Chen, Tinggui, Qianqian Li, Peihua Fu, Jianjun Yang, Chonghuan Xu, Guodong Cong, and Gongfa Li. “Public Opinion Polarization by Individual Revenue from the Social Preference Theory.” *International Journal of Environmental Research and Public Health* 17, no. 3 (February 2020). https://doi.org/10.3390/ijerph17030946.

Christin, Angèle. “From Daguerreotypes to Algorithms: Machines, Expertise, and Three Forms of Objectivity.” *ACM SIGCAS Computers and Society* 46, no. 1 (March 28, 2016): 27–32. https://doi.org/10.1145/2908216.2908220.

Chu, Derek K., Elie A. Akl, Stephanie Duda, Karla Solo, Sally Yaacoub, Holger J. Schünemann, Derek K. Chu, et al. “Physical Distancing, Face Masks, and Eye Protection to Prevent Person-to-Person Transmission of SARS-CoV-2 and COVID-19: A Systematic Review and Meta-Analysis.” *The Lancet* 395, no. 10242 (June 27, 2020): 1973–87. https://doi.org/10.1016/S0140-6736(20)31142-9.

Coleman, E. Gabriella. *Hacker, Hoaxer, Whistleblower, Spy: The Many Faces of Anonymous*. London; New York: Verso, 2014.

Correll, Michael. "Ethical Dimensions of Visualization Research." In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, 1–13. Glasgow, Scotland Uk: ACM Press, 2019. https://doi.org/10.1145/3290605.3300418.

Costanza-Chock, Sasha. *Design Justice: Towards an Intersectional Feminist Framework for Design Theory and Practice*. Cambridge, MA: MIT Press, 2020. https://papers.ssrn.com/abstract=3189696.

CrowdTangle Team. *CrowdTangle*. Menlo Park, California, United States: Facebook, 2020. https://apps.crowdtangle.com/crystalcovidvis/lists/pages.

Data for Black Lives. "Data for Black Lives COVID-19 Movement Pulse Check April 2020," April 2020.

D'Ignazio, Catherine, and Lauren F. Klein. *Data Feminism*. Strong Ideas Series. Cambridge, Massachusetts: The MIT Press, 2020.

Documenting the Now. *Twarc*. GitHub Repository, 2016. https://github.com/DocNow/twarc.

Downey, John, and Natalie Fenton. "New Media, Counter Publicity and the Public Sphere." *New Media & Society* 5, no. 2 (June 1, 2003): 185–202. https://doi.org/10.1177/1461444803005002003.

Dumit, Joseph, and Regula Valérie Burri. "Social Studies of Scientific Imaging and Visualization." In *The Handbook of Science and Technology Studies*, edited by Edward J. Hackett, Olga Amsterdamska, Judy Wajcman, and Michael Lynch, 297–317. Cambridge, MA: MIT Press, 2008.

Elisha, Omri. *Moral Ambition: Mobilization and Social Outreach in Evangelical Megachurches*. Anthropology of Christianity 12. Berkeley, CA: University of California Press, 2011.

Emerson, Robert M., Rachel I. Fretz, and Linda L. Shaw. *Writing Ethnographic Fieldnotes*. 2nd ed. Chicago Guides to Writing, Editing, and Publishing. Chicago: The University of Chicago Press, 2011.

Fauci, Anthony. "Dr. Anthony Fauci: 'Science Is Truth.'" Podcast. *Learning Curve, Department of Health and Human Services*, June 16, 2020. https://www.hhs.gov/podcasts/learning-curve/learning-curve-05-dr-anthony-fauci-science-is-truth.html.

Fischer, Frank. “Knowledge Politics and Post-Truth in Climate Denial: On the Social Construction of Alternative Facts.” *Critical Policy Studies* 13, no. 2 (April 3, 2019): 133–52. https://doi.org/10.1080/19460171.2019.1602067.

Fontichiaro, Kirstin, and Jo Angela Oehrli. “Why Data Literacy Matters.” *Knowledge Quest* 44, no. 5 (2016): 21–27.

Forrest, Jason. “How John Burn-Murdoch’s Influential Dataviz Helped the World Understand Coronavirus.” Medium. *Nightingale* (blog), April 15, 2020. https://medium.com/nightingale/how-john-burn-murdochs-influential-dataviz-helped-the-world-understand-coronavirus-6cb4a09795ae.

Geertz, Clifford. “Deep Hanging Out.” *New York Review of Books*, October 22, 1998. https://www.nybooks.com/articles/1998/10/22/deep-hanging-out/.

Gitelman, Lisa, ed. *“Raw Data” Is an Oxymoron*. Infrastructure. Cambridge, MA: MIT Press, 2013.

Guess, Andrew M., Michael Lerner, Benjamin Lyons, Jacob M. Montgomery, Brendan Nyhan, Jason Reifler, and Neelanjan Sircar. “A Digital Media Literacy Intervention Increases Discernment between Mainstream and False News in the United States and India.” *Proceedings of the National Academy of Sciences* 117, no. 27 (July 7, 2020): 15536–45. https://doi.org/10.1073/pnas.1920498117.

Hanson, Susan. “Geography and Feminism: Worlds in Collision?” *Annals of the Association of American Geographers* 82, no. 4 (1992): 569–86. https://doi.org/10.1111/j.1467-8306.1992.tb01718.x.

Haraway, Donna. “Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective.” *Feminist Studies* 14, no. 3 (1988): 575. https://doi.org/10.2307/3178066.

Harley, J B. “Deconstructing the Map.” *Cartographica: The International Journal for Geographic Information and Geovisualization* 26, no. 2 (June 1, 1989): 1–20. https://doi.org/10.3138/E635-7827-1757-9T53.

Hobbs, Renee, ed. *Exploring the Roots of Digital and Media Literacy through Personal Narrative*. Temple University Press, 2016. https://doi.org/10.2307/j.ctvrf898z.

Hochschild, Arlie Russell. *Strangers in Their Own Land: Anger and Mourning on the American Right*. New York: New Press, 2016.

Hofstadter, Richard. *Anti-Intellectualism in American Life*. 1st edition. New York: Vintage, 1966.

Hullman, Jessica. “Why Authors Don’t Visualize Uncertainty.” *IEEE Transactions on Visualization and Computer Graphics* 26, no. 1 (January 2020): 130–39. https://doi.org/10.1109/TVCG.2019.2934287.

Hullman, Jessica, and Nicholas Diakopoulos. “Visualization Rhetoric: Framing Effects in Narrative Visualization.” *IEEE Transactions on Visualization and Computer Graphics* 17, no. 12 (December 2011): 2231–40. https://doi.org/10.1109/TVCG.2011.255.

Hullman, Jessica, Xiaoli Qiao, Michael Correll, Alex Kale, and Matthew Kay. “In Pursuit of Error: A Survey of Uncertainty Visualization Evaluation.” *IEEE Transactions on Visualization and Computer Graphics* 25, no. 1 (January 2019): 903–13. https://doi.org/10.1109/TVCG.2018.2864889.

Kelly, Meghan. “Mapping Syrian Refugee Border Crossings: A Feminist Approach.” *Cartographic Perspectives*, no. 93 (November 5, 2019): 34-64-34–64. https://doi.org/10.14714/CP93.1406.

Kiesling, Scott F., Umashanthi Pavalanathan, Jim Fitzpatrick, Xiaochuang Han, and Jacob Eisenstein. “Interactional Stancetaking in Online Forums.” *Computational Linguistics* 44, no. 4 (December 2018): 683–718. https://doi.org/10.1162/coli_a_00334.

Kitchin, Rob, Martin Dodge, and Chris Perkins. “Introductory Essay: Power and Politics of Mapping.” In *The Map Reader*, 387–94. John Wiley & Sons, Ltd, 2011. https://doi.org/10.1002/9780470979587.ch50.

Kosara, Robert. “An Empire Built On Sand: Reexamining What We Think We Know About Visualization.” In *Proceedings of the Beyond Time and Errors on Novel Evaluation Methods for Visualization - BELIV ’16*, 162–68. Baltimore, MD, USA: ACM Press, 2016. https://doi.org/10.1145/2993901.2993909.

Kozinets, Robert. “Integrating: Combining Telescopic and Microscopic Understandings.” In *Netnography: The Essential Guide to Qualitative Social Media Research*, 3rd ed. London: SAGE Publications, 2019.

Kwan, Mei-Po. “Feminist Visualization: Re-Envisioning GIS as a Method in Feminist Geographic Research.” *Annals of the Association of American Geographers* 92, no. 4 (December 1, 2002): 645–61. https://doi.org/10.1111/1467-8306.00309.

Lave, Jean, and Etienne Wenger. *Situated Learning: Legitimate Peripheral Participation*. Cambridge, UK: Cambridge University Press, 1991.

Lee, Sukwon, Sung-Hee Kim, Ya-Hsin Hung, Heidi Lam, Youn-Ah Kang, and Ji Soo Yi. “How Do People Make Sense of Unfamiliar Visualizations?: A Grounded Model of Novice’s

Information Visualization Sensemaking." *IEEE Transactions on Visualization and Computer Graphics* 22, no. 1 (January 2016): 499–508. https://doi.org/10.1109/TVCG.2015.2467195.

Lee, Sukwon, Sung-Hee Kim, and Bum Chul Kwon. "VLAT: Development of a Visualization Literacy Assessment Test." *IEEE Transactions on Visualization and Computer Graphics* 23, no. 1 (January 2017): 551–60. https://doi.org/10.1109/TVCG.2016.2598920.

Lengler, R. "Identifying the Competencies of 'Visual Literacy' - a Prerequisite for Knowledge Visualization." In *Tenth International Conference on Information Visualisation (IV'06)*, 232–36. Tenth International Conference on Information Visualisation (IV'06), 2006. https://doi.org/10.1109/IV.2006.60.

Lloyd, Zenobia. "Q&A: Michael Levitt on Why There Shouldn't Be a Lockdown, How He's Been Tracking Coronavirus." *The Stanford Daily*, August 3, 2020. https://www.stanforddaily.com/2020/08/02/qa-michael-levitt-on-why-there-shouldnt-be-a-lockdown-how-hes-been-tracking-coronavirus/.

Maltese, Adam, Dubravka Svetina, and Joseph Harsh. "Research and Teaching: Data Visualization Literacy: Investigating Data Interpretation Along the Novice-Expert Continuum." *Journal of College Science Teaching* 045, no. 01 (2015). https://doi.org/10.2505/4/jcst15_045_01_84.

Markham, Annette N. "Fieldwork in Social Media." *Qualitative Communication Research* 2, no. 4 (December 1, 2013): 434–46. https://doi.org/10.1525/qcr.2013.2.4.434.

Mazza, Ed. "Top Trump DHS Official Asks Twitter To Help Him Find Coronavirus Outbreak Map." *Huffington Post*, February 25, 2020, sec. Politics. https://www.huffpost.com/entry/ken-cuccinelli-coronavirus-map_n_5e54dd40c5b66729cf618d35.

McInnes, Leland, John Healy, and James Melville. "UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction." *ArXiv:1802.03426 [Cs, Stat]*, December 6, 2018. http://arxiv.org/abs/1802.03426.

McLafferty, Sara L. "Counting for Women." *The Professional Geographer* 47, no. 4 (November 1, 1995): 436–42. https://doi.org/10.1111/j.0033-0124.1995.436_j.x.

Meyer, Miriah, and Jason Dykes. "Criteria for Rigor in Visualization Design Study." *IEEE Transactions on Visualization and Computer Graphics* 26, no. 1 (January 2020): 87–97. https://doi.org/10.1109/TVCG.2019.2934539.

Miller, Bruce L. "Science Denial and COVID Conspiracy Theories: Potential Neurological Mechanisms and Possible Responses." *Journal of the American Medical Association* 324, no. 22 (December 8, 2020): 2255. https://doi.org/10.1001/jama.2020.21332.

Moats, David, and Erik Borra. “Quali-Quantitative Methods beyond Networks: Studying Information Diffusion on Twitter with the Modulation Sequencer.” *Big Data & Society* 5, no. 1 (June 2018): 205395171877213. https://doi.org/10.1177/2053951718772137.

Mueggler, Erik. *The Paper Road: Archive and Experience in the Botanical Exploration of West China and Tibet*. Berkeley, CA: University of California Press, 2011.

Noble, Safiya Umoja. *Algorithms of Oppression: How Search Engines Reinforce Racism*. 1 edition. New York: NYU Press, 2018.

Onuoha, Mimi. “The Library of Missing Datasets,” 2016. http://mimionuoha.com/the-library-of-missing-datasets.

Oreskes, Naomi, and Erik M. Conway. *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*. New York, NY: Bloomsbury Press, 2010.

Park, Deokgun, Steven M. Drucker, Roland Fernandez, and Niklas Elmqvist. “Atom: A Grammar for Unit Visualizations.” *IEEE Transactions on Visualization and Computer Graphics* 24, no. 12 (December 1, 2018): 3032–43. https://doi.org/10.1109/TVCG.2017.2785807.

Parsons, Paul. “Promoting Representational Fluency for Cognitive Bias Mitigation in Information Visualization.” In *Cognitive Biases in Visualizations*, edited by Geoffrey Ellis, 137–47. Cham: Springer International Publishing, 2018. https://doi.org/10.1007/978-3-319-95831-6_10.

Peck, Evan M., Sofia E. Ayuso, and Omar El-Etr. “Data Is Personal: Attitudes and Perceptions of Data Visualization in Rural Pennsylvania.” In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–12. CHI ’19. New York, NY, USA: Association for Computing Machinery, 2019. https://doi.org/10.1145/3290605.3300474.

Poco, Jorge, and Jeffrey Heer. “Reverse-Engineering Visualizations: Recovering Visual Encodings from Chart Images.” *Computer Graphics Forum* 36, no. 3 (June 2017): 353–63. https://doi.org/10.1111/cgf.13193.

Porter, Theodore M. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton, NJ: Princeton University Press, 1995.

Postill, John, and Sarah Pink. “Social Media Ethnography: The Digital Researcher in a Messy Web.” *Media International Australia* 145, no. 1 (November 1, 2012): 123–34. https://doi.org/10.1177/1329878X1214500114.

Proctor, Robert N. *Golden Holocaust: Origins of the Cigarette Catastrophe and the Case for Abolition*. Berkeley, CA: University of California Press, 2011.

Ruchikachorn, Puripant, and Klaus Mueller. “Learning Visualizations by Analogy: Promoting Visual Literacy through Visualization Morphing.” *IEEE Transactions on Visualization and Computer Graphics* 21, no. 9 (September 2015): 1028–44. https://doi.org/10.1109/TVCG.2015.2413786.

Savva, Manolis, Nicholas Kong, Arti Chhajta, Li Fei-Fei, Maneesh Agrawala, and Jeffrey Heer. “ReVision: Automated Classification, Analysis and Redesign of Chart Images.” In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology - UIST ’11*, 393. Santa Barbara, California, USA: ACM Press, 2011. https://doi.org/10.1145/2047196.2047247.

Scheufele, Dietram A., and Nicole M. Krause. “Science Audiences, Misinformation, and Fake News.” *Proceedings of the National Academy of Sciences* 116, no. 16 (April 16, 2019): 7662–69. https://doi.org/10.1073/pnas.1805871115.

Segel, E, and J Heer. “Narrative Visualization: Telling Stories with Data.” *IEEE Transactions on Visualization and Computer Graphics* 16, no. 6 (November 2010): 1139–48. https://doi.org/10.1109/TVCG.2010.179.

Shneiderman, Ben. “Data Visualization’s Breakthrough Moment in the COVID-19 Crisis.” Medium. *Nightingale* (blog), April 30, 2020. https://medium.com/nightingale/data-visualizations-breakthrough-moment-in-the-covid-19-crisis-ce46627c7db5.

Small, Mario Luis. “`How Many Cases Do I Need?’: On Science and the Logic of Case Selection in Field-Based Research.” *Ethnography* 10, no. 1 (March 1, 2009): 5–38. https://doi.org/10.1177/1466138108099586.

So, Wonyoung, and Fábio Duarte. “Cartographers of North Korea: Who Are They and What Are the Technical, Political, and Social Issues Involved in Mapping North Korea.” *Geoforum* 110 (March 1, 2020): 147–56. https://doi.org/10.1016/j.geoforum.2020.02.008.

Starbird, Kate, Ahmer Arif, and Tom Wilson. “Disinformation as Collaborative Work: Surfacing the Participatory Nature of Strategic Information Operations.” *Proceedings of the ACM on Human-Computer Interaction* 3, no. CSCW (November 7, 2019): 127:1-127:26. https://doi.org/10.1145/3359229.

Stoiber, Christina, Florian Grassinger, Margit Pohl, Holger Stitz, Marc Streit, and Wolfgang Aigner. “Visualization Onboarding: Learning How to Read and Use Visualizations.” Preprint. Open Science Framework, August 15, 2019. https://doi.org/10.31219/osf.io/c38ab.

Tanahashi, Yuzuru, Nick Leaf, and Kwan-Liu Ma. “A Study On Designing Effective Introductory Materials for Information Visualization.” *Computer Graphics Forum* 35, no. 7 (2016): 117–26. https://doi.org/10.1111/cgf.13009.

Teng, Emma. *Taiwan's Imagined Geography: Chinese Colonial Travel Writing and Pictures, 1683-1895*. Harvard Univ Asia Center, 2004.

Tripodi, Francesca. "Searching for Alternative Facts." Media Manipulation. Data & Society Research Institute, n.d.

Tufekci, Zeynep. "Opinion | On Masks and Clinical Trials, Rand Paul's Tweeting Is Just Plain Wrong." *The New York Times*, November 6, 2020, sec. Opinion. https://www.nytimes.com/2020/11/06/opinion/sunday/coronavirus-masks.html.

Viegas, Fernanda B., and Martin Wattenberg. "Communication-Minded Visualization: A Call to Action [Technical Forum]." *IBM Systems Journal* 45, no. 4 (2006): 801–12. https://doi.org/10.1147/sj.454.0801.

Wenger, Etienne. *Communities of Practice: Learning, Meaning, and Identity*. Cambridge, UK; New York: Cambridge University Press, 1998.

Wu, Angela Xiao, and Harsh Taneja. "Platform Enclosure of Human Behavior and Its Measurement: Using Behavioral Trace Data against Platform Episteme." *New Media & Society*, July 1, 2020, 146144482093354. https://doi.org/10.1177/1461444820933547.

Yardi, Sarita, and danah boyd. "Dynamic Debates: An Analysis of Group Polarization Over Time on Twitter." *Bulletin of Science, Technology & Society* 30, no. 5 (October 1, 2010): 316–27. https://doi.org/10.1177/0270467610380011.

Yi-Fong Su, Vincent, Yung-Feng Yen, Kuang-Yao Yang, Wei-Juin Su, Kun-Ta Chou, Yuh-Min Chen, and Diahn-Warng Perng. "Masks and Medical Care: Two Keys to Taiwan's Success in Preventing COVID-19 Spread." *Travel Medicine and Infectious Disease*, June 4, 2020, 101780. https://doi.org/10.1016/j.tmaid.2020.101780.

Zhang, Yixuan, Yifan Sun, Sumit Barua, Enrico Bertini, and Andrea G. Parker. "Mapping the Landscape of COVID-19 Crisis Visualizations." Preprint. Open Science Framework, August 25, 2020. https://doi.org/10.31219/osf.io/kd3y9.

Chapter 3

Sociotechnical Considerations for Accessible Visualization Design

Abstract

Accessibility – the process of designing for people with disabilities (PWD) — is an important but under-explored challenge in the visualization research community. Without careful attention, and if PWD are not included as equal participants throughout the process, there is a danger of perpetuating a vision-first approach to accessible design that marginalizes the lived experience of disability (e.g., by creating overly simplistic "sensory translations" that map visual to non-visual modalities in a one-to-one fashion). In this paper, we present a set of sociotechnical considerations for research in accessible visualization design, drawing on literature in disability studies, tactile information systems, and participatory methods. We identify that using state-of-the-art technologies may introduce more barriers to access than they remove, and that expectations of research novelty may not produce outcomes well-aligned with the needs of disability communities. Instead, to promote a more inclusive design process, we emphasize the importance of clearly communicating goals, following existing accessibility guidelines, and treating PWD as equal participants who are compensated for their specialized skills. To illustrate how these considerations can be applied in practice, we discuss a case study of an inclusive design workshop held in collaboration with the Perkins School for the Blind.

Publication: This chapter appears as Alan Lundgard*, Crystal Lee* (co-first author), and Arvind Satyanarayan, "Sociotechnical Considerations for Accessible Visualization Design," *IEEE Transactions on Visualization and Computer Graphics* (October 2019).

Acknowledgements: We thank the students and staff from the Perkins School for the Blind, the workshop organizers, and the MIT Assistive Technology Information Center. We also thank Michael Correll, Richard Ladner, and Meredith Ringel Morris for their feedback on earlier drafts.

1 Introduction

History provides many instructive examples of faulty information systems designed for people with disabilities (PWD). For example, early approaches to designing tactile reading systems for blind people frequently used embossed letters that attempted to translate the alphabet's visual contours directly into tactile shapes. In the 18th century, schools favored these systems because sighted teachers did not require special training to read the embossed letters, and because it allowed blind and sighted readers to read the same texts (McGinnity, Seymour-Ford, and Andries 2004).

However, drawing on his own experience with blindness, Louis Braille—the inventor of the ubiquitous reading system—knew that raised dots were easier to read than embossed letters because they provided more tactile distinction. The eventual success of Braille's system over embossed letters is a parable for solutions that are designed by and for PWD. Since PWD are better informed about their lived experiences than designers who are not disabled, they are also in the best position to design and engineer sustainable solutions. Embossed letters were ineffective because they directly translated visual into tactile systems *without* centering a blind reader's experience.

This story resonates today as technologists extol the potential for web interfaces and 3D printing to improve the accessibility of data visualizations. Since the 1990s, visualization and Human-Computer Interaction (HCI) researchers have developed tactile interfaces that claim to make information more accessible to blind people by allowing users to touch, rather than see, the data (Choi et al. 2019, Ishii and Ullmer 1997, Jansen et al. 2015). However, few of these projects extensively consult blind people, which means that they often become what disability design expert Liz Jackson calls *disability dongles*: a "well intended, elegant, yet useless solution to a problem we [PWD] never knew we had" (Jackson 2019). Developments in HCI and tactile graphics have previously converged in Assistive Technology (AT) research (Bennett, Brady, and Branham 2018; Ladner et al. 2005, Muller 2003), but visualization researchers have not yet substantively engaged with the developments in tactile information systems, well-established accessibility guidelines, or with the perspectives of PWD.

To begin addressing these issues, we present a set of *sociotechnical* considerations for accessible visualization design. In our analysis, we emphasize that there are both *social* and *technological* constraints on the accessible visualization design space. We demonstrate how our considerations—informed by work written by and for PWD in disability studies, AT, and design history—can be applied to future accessibility research. We also critique current research that emphasizes a vision-first approach to making tactile graphics.

2 Background & Related Work

Our work brings together disability studies, research on tactile information systems, and participatory design methods to articulate a set of sociotechnical considerations for accessible visualization design.

2.1 Disability Studies

In the United States, disability studies is an interdisciplinary field that seeks to untangle the different political, intellectual, and cultural dimensions of disability in society. In this literature, scholars distinguish between two models of thinking about disability (Shakespeare 2013). In a **medical model**, the disability diagnosis is tied directly to an individual's physical or psychological state, and the prognosis focuses on curing or managing the disability until it disappears as much as possible (e.g., a deaf person who uses cochlear implants; Hamraie 2012). By contrast, a **social model** of disability distinguishes between an *impairment* (a physical or psychological abnormality) and a *disability*, which describes a form of social, economic, and political exclusion perpetuated against people who have some impairment (Oliver 2013). In the social model, people are disabled by structural and environmental factors, not by their bodies (Thomson 1996).

While these two models are not mutually exclusive, and can be considered alongside personal, environmental, and other contextual factors ("World Report on Disability 2011" 2011), it is important for visualization researchers to understand the social and medical models because each can be leveraged towards different kinds of accessibility research. The medical model, for example, clearly articulates an actionable problem that can be measured and fixed, which can be practically useful to AT researchers. However, AT research that focuses exclusively on the medical model often stigmatizes PWD by attempting to "fix" all impairments and differences through medical intervention, even when the condition does not cause pain, illness, or death (e.g., deafness) (Davis 2013). In particular, the medical model misses how an impairment *develops* into a disability. For example, in a society where everyone communicates through sign language, deafness may not be perceived as a disability at all, but may instead be a source of cultural pride. In a hearing society, however, deafness may be seen as a disability that requires accommodation through interpreters or medical intervention (Mankoff, Hayes, and Kasnitz 2010; Mills 2015).

The social model of disability, by contrast, shifts the focus from "cure to care," in which the goal is not to fix the impairment, but rather the ways that society treats PWD (Zola 1983). For example, rather than building a better wheelchair that allows its user to climb stairs, proponents of the social model look instead to dismantling the legislative or bureaucratic barriers to installing wheelchair ramps in public places. While the social model cannot explain abnormal biological functions, it can be crucial for driving the adoption of assistive technologies and ensuring that the technologies address the infrastructural barriers that PWD experience in their everyday lives.

2.2 Tactile Information Systems

Tactile systems like braille displays, 3D models, and embossed maps have been a mainstay of blind education (Van Geem n.d.). As such, there is a vast literature by computer scientists, education professionals, and organizations for the blind that explore how these tools can be put into practice (Zebehazy and Wilton 2014; Gardiner and Perkins 1997; Schneider 2000). Education researcher Lucia Hasty, for example, has developed principles of graphical literacy that show how blind and sighted students learn differently (Hasty 2017). Sighted learners absorb information **whole to part**, where they see the whole picture simultaneously, and understand the different visual encodings in relation to each other (e.g., the size and color of one bar compared to another). Tactile learners, however, must approach tactile graphics **part to whole** by touching individual parts of the graphic processually. They then put together each piece of information in a sequence to view the graphic as a whole.

Given the differences between these approaches, successful tactile representations should ensure that it is easy to access the information sequentially, and that each tactile element is easily and quickly distinguishable. Braille is a prime example of a successful tactile representation. Each braille cell is easy to scan quickly and sequentially (compared to embossed lettering, where a user must try to synthesize each letter individually—a relatively slow process—before moving onto the next letter). Furthermore, printed braille has many well-established contractions, which make it even faster to scan through texts. Drawing on braille's design successes and guidelines, an important area for future work is developing an analogous system for visualizations (i.e., one that permits fast and effective information access, and does not simply translate the visualization into a tactile representation in a one-to-one fashion). Combining these insights with guidelines on tactile graphics from organizations like American Printing House for the Blind (APH), Braille Authority of North America (BANA), and World Wide Web Consortium (W3C)'s Web Accessibility Initiative (WAI), provide a powerful basis on which to begin this research (Braille Authority of North America 2010, "Web Content Accessibility Guidelines (WCAG)" 2018; Amick and Corcoran 1997). Building on these existing AT guidelines is an excellent example of how the social model of disability can inform practice (Mankoff, Hayes, and Kasnitz 2010).

2.3 Participatory Design Methods

At each stage of conducting a study, researchers should be attentive to the power dynamics between themselves and their study participants. Historically, ignoring these dynamics has exposed already marginalized communities to long-term psychological and physiological harm (e.g., the Tuskegee syphilis experiment, Brandt 1978). In particular, developments in AT have often relied upon using PWD as test subjects for technologies that were later transformed into more profitable ventures intended for able-bodied people [29]. To mitigate situations like these, researchers in many disciplines have developed ways of working *with* (rather than *on*) marginalized communities (Cargo and Mercer 2008; Costanza-Chock 2020, Moree 2018). Generally, these community-based or participatory research methods emphasize collective inquiry in which study participants are considered co-researchers. Far from simply being a *subject* of research, these participants help define the design problem and contribute to methods, data collection, analysis, and publication (Mankoff, Hayes, and Kasnitz 2010; Muller 2003; Muller and Kuhn 1993). These methods go beyond user-centered design (Norman and Draper 1986) to articulate more clearly what the stakes are in a research project: who participates with whom in what? Who are the intended beneficiaries of a project, and how do they accrue these benefits? By re-centering PWD in the design process, participatory methods can be useful for historically marginalized communities because they break down the separation between those who are doing the design and those who are being designed *for*.

This work also attempts to avoid the problems associated with *parachute research*, a phenomenon in which researchers— particularly those from wealthy universities—drop into a community, make use of local infrastructure and expertise, and then disengage from the community altogether after publishing results in a prestigious academic journal [20]. This kind of research is harmful to the communities who take the time and resources to help facilitate academic research without reciprocal benefits. This can have a disproportionately negative impact on disability communities, who already face many barriers to participating in public life. To address these problems, scholars in AT, disability studies, and design have emphasized how researchers need to consider their participants as active agents with ideas and goals that may conflict with those of the researcher. This should not be seen as an obstacle to research; rather, it provides new opportunities for collaborative design (Howard and Irani 2019; Riles 2015).

3 Sociotechnical Considerations

In this section, we introduce a set of sociotechnical considerations for accessible visualization design. These considerations are informed by related work, our own decade-long engagement with the blind community, and a case study of an inclusive design workshop. Each consideration begins by explaining its social and technological aspects, and concludes with questions that researchers and designers should consider when collaborating with disability communities.

Non-Intervention. At each step of the design process, researchers should consider whether any technological intervention is appropriate at all (Baumer and Silberman 2011). This is especially true of designing for and with PWD because well-meaning interventions may worsen the very situations they are intended to help (Jackson 2019). Design processes are actively harmful when they exhaust collaborators' time and resources without adequate compensation and reciprocity, or when they are simply band-aid technological solutions to infrastructural problems. Is there an equally viable low-tech or no-tech solution? Might the technological intervention result in more harm than the problem it is meant to address? Does the technology solve a computationally tractable version of the problem, or does it address an actual need?

Research & Design. If technological intervention is appropriate, researchers should carefully evaluate the goals of their project. In HCI, there is a well-known tension between doing "research" and doing "design" (Olson and Kellogg 2014; Zimmerman, Forlizzi, and Evenson 2007). On the one hand, research has the goal of creating new knowledge, but this may not be the best solution for a user's immediate needs. On the other hand, design should satisfy those needs, but the solution may lack research novelty. The two are not mutually exclusive, but they do prescribe different methods and goals (e.g., a research publication versus an adoptable design solution). This tension is especially apparent when designing for and with PWD. Preliminary interviews may reveal that the best design solutions to a user's needs will not count as publication-worthy material. Is the proposed solution addressing those needs, or is it a novel research contribution? Is there a solution that addresses both?

Participatory Methods. Taking a participatory approach to research and design is a key tenet of accessibility. This involves including stakeholders in the design process from the outset, during need-finding, problem definition, prototyping, publication, and dissemination (Mankoff, Hayes, and Kasnitz 2010). Participatory design aims to include users as equal participants in the design process, as opposed to merely verifying the usefulness of a solution via user studies (Dourish 2006). Inclusive design, as complementary to participatory design, places emphasis on empowering as many people as possible, often with a focus on removing barriers for people who have physical or cognitive disabilities (Muller and Kuhn 1993). Inclusive approaches may involve consideration of the *interdependence* between researchers and stakeholders (i.e., collapsing the distinction between who is doing the research, and *for whom* that research is being done) (Bennett, Brady, and Branham 2018; Bennett and Rosner 2019). Who are the stakeholders in a research project? Are they all recurring and equal participants throughout the design process? Are the research and design tools themselves accessible (Ladner 2015)?

Communicating Expectations. Because HCI and AT research projects can often be one-off or proof-of-concept prototypes that are no longer maintained after publication, it is especially important for researchers to communicate their intentions, expectations, and capabilities clearly to all collaborators. This may include the intended duration of the project, the quantity of available resources, and the project goal—whether it be a research publication or adoptable design solution. This is especially important for technical projects that require maintenance or engineering support beyond the duration of the project. How long will a project be maintained, and what are the expected contributions of each collaborator? If the goal is academic publication, has there been a discussion about author credit and order? If the goal is a marketable prototype, has there been a discussion about equitable compensation and intellectual property?

Time & Compensation. It is generally good practice to be sensitive and respectful of all collaborators' time, but this is especially true when designing for and with PWD. Like everyone, PWD are busy, but they must also contend with additional barriers to mobility, access, and employment. This makes the time that they spend on a design project especially valuable. As with any specialist, PWD should always be compensated at a rate commensurate with their specialized skills, such as the ability to read braille or use a screen reader. Additionally, it may be difficult to find and recruit PWD for collaboration, and this should be reflected in their rate of compensation, meaning a rate that is greater than that of an average user study participant (e.g., in some cases $35.00 per hour) (Bigham, Lin, and Savage 2017; Oney et al. 2018). How much time does each participant have to contribute to the collaboration? Would it be better to not participate if it is not possible to contribute meaningfully or to bring the project to fruition? Are participants with disabilities being compensated adequately and fairly?

Accessibility Guidelines. Researchers should familiarize themselves with the accessibility guidelines that have been developed by major organizations like the APH, BANA, and the W3C's WAI (Braille Authority of North America 2010; "Web Content Accessibility Guidelines (WCAG)" 2018; Amick and Corcoran 1997; Gardiner and Perkins 1997). While these guidelines are not complete and are sometimes contested, researchers should use them to steer and expand their inquiry. For example, accessibility guidelines for 3D printed tactile graphics might build upon the BANA guidelines, and accessible web-based visualizations might integrate Accessible Rich Internet Application (ARIA) attributes developed by the W3C. Are there existing standards and best practices relevant to the design problem? If so, does the proposed design solution adhere to these standards? If not, in what way might the proposed design solution integrate with, or build upon, this work?

Technology Access. Designing for and with PWD almost always involves securing access to specialized materials and technologies, from low-tech solutions (e.g., tactile tape, puff paint, wax strips) to high-tech devices (e.g., screen readers, refreshable braille displays, embossing and 3D printers). Access to these technologies will circumscribe the space of feasible design solutions, and so particular approaches should be chosen carefully to fit within these constraints. A high-tech approach may permit finer-grained, more durable products, and it may yield more sensational results, but securing access to these technologies can be prohibitively costly. A low-tech approach may be more readily available to the user on a daily basis, but may convey information more coarsely or unreliably. Even ostensibly democratizing technologies, such as 3D printing, may be hard to come by, and the interfaces to operate these technologies may not be accessible themselves. Does the problem require a high-tech approach, or will a low-tech approach work just as well, if not better? Does the design require one-time access to an expensive piece of technology, or repeated, frequent access by the user?

Technology Resolution. In addition to constraints imposed by access to specialized technologies, researchers also need to ensure that those technologies can encode information effectively. By drawing an analogy with high-fidelity audio and display resolution, we use "resolution" to describe how well a particular medium can encode and convey detailed information. This is especially important for web visualizations, which are typically not screen reader compatible, and for braille, which requires standardized height and spacing, and cannot be resized to fit a particular area. For example, with 3D printing, the size of the printer bed may constrain the amount of braille that can be printed in one line, and the quality of the printing material may change how long a braille reader can interact with the object. Plastic filaments used on consumer 3D printers, for example, can create abrasive surfaces that make it uncomfortable for braille readers to use for a long period of time. A medium's resolution depends on various hardware and software limitations, and on material and social constraints. What degree of resolution is appropriate to the user's needs, and which technologies can be used to achieve it?

4 Case Study: The Perkins School for the Blind

To illustrate the benefits of approaching accessible visualization design from a disability studies and AT perspective, and to ground our sociotechnical considerations in a concrete example, we present a case study from our own work. We emphasize that this is not an ideal case, but a *problem case* exemplifying some of the pitfalls articulated in our set of considerations. Throughout this section, we refer back to these considerations in parentheses.

The authors took part in an inter-semester inclusive design workshop that featured a collaboration between MIT and the Perkins School for the Blind. The express agreement of this collaboration was to make design interventions that would address the needs of the Perkins School students and staff **(Non-Intervention).** While the Perkins School had previously collaborated with other institutions and entrepreneurs, our collaboration involved multiple need-finding visits to the school's campus during which we toured the facilities and participated in discussions with occupational therapists, teachers, and technologists. This afforded opportunities for design iteration and feedback **(Participatory Methods).** Of prior non-participatory collaborations, the President of the Perkins School noted that:

> "These entrepreneurs come to us with their finished prototypes, but they haven't talked to very many blind people in the process before they've put in all the effort to create the prototype."

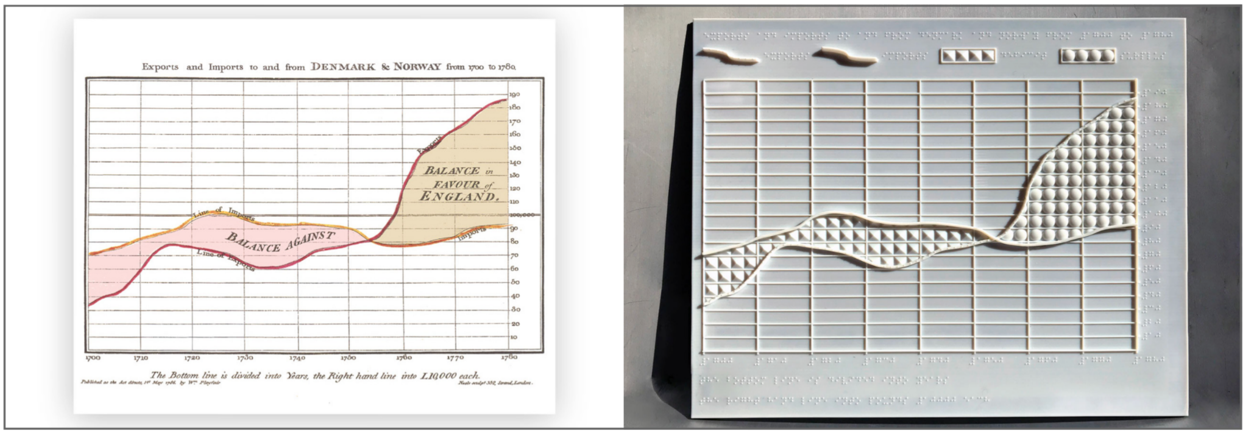

Figure 1. As part of an inclusive design workshop at the Perkins School for the Blind, we created a 3D printed tactile translation of a time-series chart by William Playfair. In this paper, we show how these one-to-one translations, while based on existing best-practice guidelines for tactile graphics, can be pedagogically ineffective and incur prohibitive costs.

In such cases, the President suggested that many of the resulting prototypes were relatively useless to their community, reflecting a sighted designer's idea of what a blind person *could* want, and not what they *actually* needed. By taking a participatory approach, the workshop was meant to evade this pitfall. However, we identified an ongoing tension between doing "research" and doing "design" which made clear that participatory methods are necessary but insufficient for framing design interventions. Although we did not enter the workshop intending to produce academic research, the workshop organizers placed an emphasis on developing a novel research product. They encouraged us to "think beyond" satisfying the commonplace and bespoke needs of the Perkins School students and staff towards new, more generalizable solutions **(Research** & **Design).** To relieve this tension, we opted to work on a project that engaged with the immediate needs of blind students, but also had potential for research novelty—3D printed tactile graphics. One assistive technologist, for example, noted that prior approaches to tactile graphics in educational settings were ineffective:

> “I have not found a tactile graphic solution that works. There's usually some sort of compromise. You might get the general outline of something, but you don't necessarily get the information that is conveyed… Now, if there was a new student, I would turn to 3D printing to create 3D visualizations.”

Because tactile graphics are often created using embossing printers that are limited in the number of printable dots-per-inch, direct translation of a high resolution visual image to a lower resolution tactile print leads to information loss. This can create tactile ambiguities that were not present in the original visual image. For example, in a paper embossed tactile graphic, the intersections between two lines in a multi-line plot would not be distinguishable through touch, as in the original Playfair chart (Figure 1). 3D printing permits higher resolution encodings because there is an extra printable dimension, which helps eliminate the ambiguities common to paper embossing **(Technology Resolution).** Accordingly, 3D printing is often touted as a democratizing technology in terms of affordability and access, but the more affordable models (under $800) also have the lowest resolution and smallest printable dimensions. These constraints on print resolution may be incompatible with constraints imposed by existing accessibility guidelines. For example, BANA gives guidelines for the spacing between braille dots because braille becomes illegible when it is spaced too widely or narrowly (Braille Authority of North America 2010). Due to their low resolution and small printable dimensions, prototypes of the Playfair translation printed using the affordable 3D printers failed to meet the BANA guidelines **(Accessibility Guidelines).** Thus, access to an extremely expensive commercial 3D printer (roughly $330,000) became a prerequisite for translating the detailed visualization into a high-resolution tactile graphic **(Technology Access).** Compared with paper embossed graphics, the Perkins School students expressed a strong preference for the 3D printed graphics because of their higher resolution. This is a worthwhile research insight, but commercial 3D printing was not a design solution that was adoptable beyond the duration of the workshop and should have been communicated as such **(Communicating Expectations).**

The collaboration between MIT and the Perkins School was well-meaning and it was intended to be mutually beneficial. As such, the Perkins School students and staff were not compensated for their valuable time and expertise, even though they spent part of their workdays to meet with us **(Time & Compensation)**. While this arrangement may have been equitable had we collaboratively produced solutions that actually satisfied the needs of the students, it was not clear that the workshop benefited the Perkins School students and staff as much as it benefited the participants from MIT. Media scholar Mara Mills has documented the many ways that disability has been used as a pretense to develop innovations that are primarily for publicity, often without giving back to PWD in a substantial way (Mills 2010). While this was in no way the workshop's intent, it may have been its predominant outcome: an opportunity for able-bodied researchers and designers to engage with PWD and to generate publicity for both institutions (both of which had reporters on-hand). For us, however, the workshop also afforded a valuable lesson for guiding future research and design that avoids the pitfalls of parachute research. Put succinctly, successful participatory design cannot be achieved within an accelerated time frame. Collaborations with PWD should support longer-term engagement through equitable compensation for each participant's time, and the goals of the design process should be well-scoped to account for each participant's availability and access to relevant technologies.

5 Conclusion & Future Work

In this paper, we contribute a set of sociotechnical considerations for accessible visualization design. Visualization research has largely focused on addressing the technological and perceptual barriers to the effective visual encoding of information. However, due to the many ways in which barriers to access are manifested, we em-phasize that there are both *social* and *technological* constraints on the accessible visualization design space. While new technologies such as 3D printing and web interfaces afford many opportunities for future work (e.g., developing more effective tactile representations, guidelines for 3D printing, and screen reader compatible visualizations), existing work on braille and tactile graphics provide a guide for conducting this research successfully and inclusively. Drawing on these examples, our considerations are intended to help guide researchers away from *parachute research*, overly simplistic vision-first approaches (e.g., *disability dongles*), and towards design practices that avoid the pitfalls of well-meaning but insufficient collaborations with PWD (e.g., embossed lettering). When pursued with careful attention, there will be many exciting opportunities for collaboration between disability, AT, and visualization communities.

References for Chapter 3

Aldrich, Frances. "Talk to the Hand: An Agenda for Further Research on Tactile Graphics." In *Diagrammatic Representation and Inference*, edited by Gem Stapleton, John Howse, and John Lee, 344–46. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2008.

Amick, Nancy, and Jane Corcoran. "Guidelines for Design of Tactile Graphics." American Printing House for the Blind, 1997. https://www.aph.org/research/guides/.

Baumer, Eric P.S., and M. Six Silberman. "When the Implication Is Not to Design (Technology)." In *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems - CHI '11*, 2271. Vancouver, BC, Canada: ACM Press, 2011. https://doi.org/10.1145/1978942.1979275.

Bennett, Cynthia L., 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*, 161–73. Galway, Ireland: ACM Press, 2018. https://doi.org/10.1145/3234695.3236348.

Bigham, Jeffrey P., Irene Lin, and Saiph Savage. "The Effects of 'Not Knowing What You Don't Know' on Web Accessibility for Blind Web Users." In *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility*, 101–9. Baltimore Maryland USA: ACM, 2017. https://doi.org/10.1145/3132525.3132533.

Braille Authority of North America. "Guidelines and Standards for Tactile Graphics, 2010." Braille Authority of North America, 2010. http://www.brailleauthority.org/tg/web-manual/index.html.

Brandt, Allan M. "Racism and Research: The Case of the Tuskegee Syphilis Study." *The Hastings Center Report* 8, no. 6 (1978): 21–29. https://doi.org/10.2307/3561468.

Cargo, Margaret, and Shawna L. Mercer. "The Value and Challenges of Participatory Research: Strengthening Its Practice." *Annual Review of Public Health* 29, no. 1 (2008): 325–50. https://doi.org/10.1146/annurev.publhealth.29.091307.083824.

Choi, Jinho, Sanghun Jung, Deok Gun Park, Jaegul Choo, and Niklas Elmqvist. "Visualizing for the Non-Visual: Enabling the Visually Impaired to Use Visualization." *Computer Graphics Forum* 38, no. 3 (June 2019): 249–60. https://doi.org/10.1111/cgf.13686.

Costanza-Chock, Sasha. *Design Justice: Towards an Intersectional Feminist Framework for Design Theory and Practice*. Cambridge, MA: MIT Press, 2020. https://papers.ssrn.com/abstract=3189696.

Davis, Lennard J. *The Disability Studies Reader*. 4th ed. New York, NY: Routledge, 2013.

Dourish, Paul. “Implications for Design.” In *CHI Conference on Human Factors in Computing Systems*, 541–50. Montreal, Canada, 2006.

Gardiner, Ann, and Chris Perkins. “Best Practice Guidelines for the Design, Production and Presentation of Vacuum Formed Tactile Maps.” Tactile Book Advancement Group, 1997. http://www.tactilebooks.org/tactileguidelines/page1.htm.

Hamraie, Aimi. “Universal Design Research as a New Materialist Practice.” *Disability Studies Quarterly* 32, no. 4 (2012).

Hanson, Vicki L., Anna Cavender, and Shari Trewin. “Writing about Accessibility.” *Interactions* 22, no. 6 (October 28, 2015): 62–65. https://doi.org/10.1145/2828432.

Hasty, Lucia. “Teaching Tactile Graphics.” Perkins School for the Blind. Perkins School for the Blind, 2017. http://www.perkinselearning.org/videos/webcast/teaching-tactile-graphics.

Health, The Lancet Global. “Closing the Door on Parachutes and Parasites.” *The Lancet Global Health* 6, no. 6 (June 1, 2018): e593. https://doi.org/10.1016/S2214-109X(18)30239-0.

Howard, Dorothy, and Lilly Irani. “Ways of Knowing When Research Subjects Care.” In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI ’19*, 1–16. Glasgow, Scotland: ACM Press, 2019. https://doi.org/10.1145/3290605.3300327.

Ishii, Hiroshi, and Brygg Ullmer. “Tangible Bits: Towards Seamless Interfaces Between People, Bits and Atoms.” In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, 234–41. CHI ’97. New York, NY, USA: ACM, 1997. https://doi.org/10.1145/258549.258715.

Jackson, Liz. “Disability Dongle: A Well-Intended, Elegant, yet Useless Solution to a Problem We Never Knew We Had. Disability Dongles Are Most Often Conceived of and Created in Design Schools and at IDEO.” Tweet. Twitter, March 2019. https://twitter.com/elizejackson/status/1110629818234818570.

Jansen, Yvonne, Pierre Dragicevic, Petra Isenberg, Jason Alexander, Abhijit Karnik, Johan Kildal, Sriram Subramanian, and Kasper Hornbæk. “Opportunities and Challenges for Data Physicalization.” In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI ’15*, 3227–36. Seoul, Republic of Korea: ACM Press, 2015. https://doi.org/10.1145/2702123.2702180.

Ladner, Richard E. “Design for User Empowerment.” *Interactions* 22, no. 2 (February 2015): 24–29. https://doi.org/10.1145/2723869.

Ladner, Richard E., Beverly Slabosky, Andrew Martin, Amelia Lacenski, Stuart Olsen, Dmitri Groce, Melody Y. Ivory, et al. "Automating Tactile Graphics Translation." In *Proceedings of the 7th International ACM SIGACCESS Conference on Computers and Accessibility - Assets '05*, 150. Baltimore, MD, USA: ACM Press, 2005. https://doi.org/10.1145/1090785.1090814.

Mankoff, Jennifer, Gillian R. Hayes, and Devva Kasnitz. "Disability Studies as a Source of Critical Inquiry for the Field of Assistive Technology." In *Proceedings of the 12th International ACM SIGACCESS Conference on Computers and Accessibility - ASSETS '10*, 3. Orlando, Florida, USA: ACM Press, 2010. https://doi.org/10.1145/1878803.1878807.

McGinnity, B. L., J. Seymour-Ford, and K. J. Andries. "Books for the Blind." Perkins History Museum, Perkins School for the Blind, 2004. https://www.perkins.org/history/legacy/books-for-the-blind.

Mills, Mara. "Deaf Jam: From Inscription to Reproduction to Information." *Social Text* 28, no. 102 (March 2010): 35–58. https://doi.org/10.1215/01642472-2009-059.

———. "Deafness." In *Keywords in Sound*, edited by David Novak and Matt Sakakeeny. Durham, NC: Duke University Press, 2015.

Moree, Dana. "Qualitative Approaches to Studying Marginalized Communities." *Oxford Research Encyclopedia of Education*, July 30, 2018. https://doi.org/10.1093/acrefore/9780190264093.013.246.

Muller, Michael J. "Participatory Design: The Third Space in HCI." In *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications, Second Edition*, edited by Andrew Sears and Julie A. Jacko, 1061–81. CRC Press, 2007. https://doi.org/10.1201/9781410615862.

Muller, Michael J., and Sarah Kuhn. "Participatory Design." *Communications of the ACM* 36, no. 6 (June 1, 1993): 24–28. https://doi.org/10.1145/153571.255960.

Norman, Donald A., and Stephen W. Draper. *User Centered System Design; New Perspectives on Human-Computer Interaction*. Hillsdale, NJ, USA: L. Erlbaum Associates Inc., 1986.

Oliver, Mike. "The Social Model of Disability: Thirty Years On." *Disability & Society* 28, no. 7 (2013): 1024–26.

Olson, Judith S., and Wendy Kellogg, eds. *Ways of Knowing in HCI*. New York: Springer, 2014.

Oney, Steve, Alan Lundgard, Rebecca Krosnick, Michael Nebeling, and Walter S. Lasecki. "Arboretum and Arbility: Improving Web Accessibility Through a Shared Browsing Architecture." In *Proceedings of the 31st Annual ACM Symposium on User Interface*

Software and Technology, 937–49. Berlin Germany: ACM, 2018. https://doi.org/10.1145/3242587.3242649.

Riles, Annelise. “From Comparison to Collaboration: Experiments with a New Scholarly and Political Form.” *Law and Contemporary Problems* 78 (2015): 147–83.

Schneider, Jochen. “Constructing the Yellow Brick Road: Route Bricks on Virtual Tactile Maps.” In *Proceedings of the International Conference on Computers Helping People with Special Needs (ICCHP)*, 641–48. Karlsruhe, Germany, 2000.

Shakespeare, Tom. “The Social Model of Disability.” In *The Disability Studies Reader*, edited by Lennard J. Davis, 4th ed., 197–204. New York, NY: Routledge, 2013.

Thomson, Rosemarie Garland. *Extraordinary Bodies: Figuring Physical Disability in American Culture and Literature*. 1st edition. New York: Columbia University Press, 1996.

Van Geem, Patrick. “Tactile Graphics Resources.” Texas School for the Blind and Visually Impaired. Accessed June 12, 2019. https://www.tsbvi.edu/math/107-graphics/3189-tactile-graphics-resources.

W3C Web Accessibility Initiative (WAI). “Web Content Accessibility Guidelines (WCAG),” June 22, 2018. https://www.w3.org/WAI/standards-guidelines/wcag/.

“World Report on Disability 2011.” Geneva, Switzerland: World Health Organization, 2011. http://www.ncbi.nlm.nih.gov/books/NBK304079/.

Zimmerman, John, 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*, 493–502. San Jose California USA: ACM, 2007. https://doi.org/10.1145/1240624.1240704.

Zola, Irving Kenneth. *Socio-Medical Inquiries: Recollections, Reflections and Reconsiderations*. Philadelphia: Temple University Press, 1983.

Chapter 4

Rich Screen Reader Experiences for Accessible Data Visualization

Abstract

Current web accessibility guidelines ask visualization designers to support screen readers via basic non-visual alternatives like textual descriptions and access to raw data tables. But charts do more than summarize data or reproduce tables; they afford interactive data exploration at varying levels of granularity — from fine-grained datum-by-datum reading to skimming and surfacing high-level trends. In response to the lack of comparable non-visual affordances, we present a set of rich screen reader experiences for accessible data visualization and exploration. Through an iterative co-design process, we identify three key design dimensions for expressive screen reader accessibility: structure, or how chart entities should be organized for a screen reader to traverse; navigation, or the structural, spatial, and targeted operations a user might perform to step through the structure; and, description, or the semantic content, composition, and verbosity of the screen reader's narration. We operationalize these dimensions to prototype screen-reader-accessible visualizations that cover a diverse range of chart types and combinations of our design dimensions. We evaluate a subset of these prototypes in a mixed-methods study with 13 blind and visually impaired readers. Our findings demonstrate that these designs help users conceptualize data spatially, selectively attend to data of interest at different levels of granularity, and experience control and agency over their data analysis process.

Publication: This chapter appears as Jonathan Zong*, Crystal Lee*, Alan Lundgard*, JiWoong Jang, Daniel Hajas, and Arvind Satyanarayan, "Sociotechnical Considerations for Accessible Visualization Design," *IEEE Transactions on Visualization and Computer Graphics* (October 2019).

Acknowledgements: We thank Matthew Blanco, Evan Peck, the NYU Digital Theory Lab, and the MIT Accessibility Office. This work was supported by NSF awards #1942659, #1122374, and #1941577.

1 Introduction

Despite decades of visualization research and recent legal requirements to make web-based content accessible (W3C 2018; Higgins 2019), web-based visualizations remain largely inaccessible to people with visual disabilities. Charts on mainstream publications are often completely invisible to screen readers (an assistive technology that transforms text and visual media into speech) or are rendered as incomprehensible strings of “graphic graphic graphic” (Sharif et al. 2021; Fossheim 2020). Current accessibility guidelines ask visualization designers to provide textual descriptions of their graphics via alt text (short for alternative text) and link to underlying data tables (Gould, O’Connell, and Freed 2008; W3C 2019). However, these recommendations do not provide modes of information-seeking comparable to what sighted readers enjoy with interactive visualizations. For instance, well-written alt text can provide a high-level takeaway of what the visualization shows, but it does not allow readers to drill down into the data to explore specific sections. While tables provide readers with the ability to hone in on specific data points, reading data line-by-line quickly becomes tedious and makes it difficult to identify overall trends.

Developing rich non-visual screen reader experiences for data visualizations poses several unique challenges. Although visuomotor interactions (like hovering, pointing, clicking, and dragging) have been core to visualization research (Dimara and Perin 2020), screen readers redefine what interaction is for visualization. Rather than primarily manipulating aspects of the visualization or its backing data pipeline (Yi et al. 2007; Heer and Shneiderman 2012; Dimara and Perin 2020), screen readers make reading a visualization an interactive operation as well — users must intentionally perform actions with their input devices in order to cognize visualized elements. Moreover, as screen readers narrate elements one-at-a-time, they explicitly linearize reading a visualization. As a result, in contrast to sighted readers who can choose to selectively attend to specific elements and have access to the entire visualization during the reading process, screen reader users are limited to the linear steps made available by the visualization author and must remember (or note down) prior output conveyed by the screen reader. Despite these modality differences, studies have found that screen reader users share the same information-seeking goals as sighted readers: an initial holistic overview followed by comparing data points (Sharif et al. 2021), akin to the information-seeking mantra of “overview first, zoom and filter, and details on demand” (Shneiderman 2003).

In this paper, we begin to bridge this divide by conducting an iterative co-design process (co-author Hajas is a blind researcher with relevant experience) prototyping rich and usable screen reader experiences for web-based visualizations. We identify three design dimensions for enabling an expressive space of experiences: structure, or how the different elements of a chart should be organized for a screen reader to traverse; navigation, which describes the operations a user may perform to move through this structure; and, description, which specifies the semantic content, composition, and verbosity of text conveyed at each step. We demonstrate how to operationalize these design dimensions through diverse accessible reading experiences across a variety of chart types.

To evaluate our contribution, we conduct an exploratory mixed-methods study with a subset of our prototypes and 13 blind or visually impaired screen reader users. We identify specific features that make visualizations more useful for screen reader users (e.g., hierarchical and segmented approaches to presenting data, cursors and road maps for spatial navigation) and identify behavior patterns that screen reader users follow as they read a visualization (e.g., constant hypothesis testing and validating their mental models).

2 Background and Related Work

Screen Reader Assistive Technology. A screen reader is an assistive technology that conveys digital text or images as synthesized speech or braille output. Screen readers are available as standalone third-party software or can be built-in features of desktop and mobile operating systems. A screen reader allows a user to navigate content linearly with input methods native to a given platform (e.g., touch on smartphones, mouse/keyboard input on desktop). Content authors must generate and attach alt text to their visual content like images or charts in order for them to be accessible to screen reader users. Functionality and user experience differs across platforms and screen readers. In this paper, however, we focus on interacting with web-based visualizations with the most widely used desktop screen readers (JAWS/NVDA for Windows, VoiceOver for Mac).

Web Accessibility Standards. In 2014, the World Wide Web Consortium (W3C) adopted the Web Accessibility Initiative's Accessible Rich Internet Applications protocol (WAI-ARIA) which introduced a range of semantically-meaningful HTML attributes to allow screen readers to better parse HTML elements (MDN Contributors 2021). In particular, these attributes allow a screen reader to convey the state of dynamic widgets (e.g., autocomplete is available for text entry), alert users to live content updates, and identify common sections of a web page for rapid navigation (e.g., banners or the main content). In 2018, the W3C published the WAI-ARIA Graphics Module (W3C 2018) with additional attributes to support marking up structured graphics such as charts, maps, and diagrams. These attributes allow designers to annotate individual and groups of graphical elements as well as surface data values and labels for a screen reader to read aloud.

Accessible Visualization Design. In a recent survey, Kim et al. (2021) describe the rich body of work that has explored multi-sensory approaches to visualization for multiple disabilities (Yu, Ramloll, and Brewster 2001; Hasty et al. 2011; Kaper, Wiebel, and Tipei 1999; Barrass and Kramer 1999; Wu et al. 2021; Lundgard, Lee, and Satyanarayan 2019). Here, we focus on screen reader output native to web-based interfaces for blind users (namely via speech). Sharif et al. (2021) find that many web-based charts are intentionally designed to cause screen readers to skip over them. For charts that a screen reader does detect, blind or visually impaired users nevertheless experience significant difficulties: these users spend 211% more time interacting with the charts and are 61% less accurate in extracting information compared to non-screen-reader users (ibid., 2). Despite the availability of ARIA, alt text and data tables remain the most commonly used and recommended methods for making web-based charts accessible to screen readers (Gould, O'Connell, and Freed 2008; W3C 2019; Choi et al. 2019). However, each of these three approaches comes with its own limitations. Static alt text requires blind readers to accept the author's interpretation of the data; by not affording exploratory and interactive modes, alt text robs readers of the necessary time and space to interpret the numbers for themselves (Lundgard and Satyanarayan 2021). Recent research also suggests that blind people have nuanced preferences for the kinds of visual semantic content conveyed via text (Potluri et al. 2021; Lundgard and Satyanarayan 2021), and desire more interactive and exploratory representations of pictorial images (Morris et al. 2018). Data tables, on the other hand, undo the benefits of abstraction that visualizations enable — they force readers to step sequentially through data values making it difficult to identify larger-scale patterns or trends, and do not leverage the structure inherent to web-based grammars of graphics (Bostock, Ogievetsky, and Heer 2011, Satyanarayan et al. 2017). Finally, ARIA labels are not a panacea; even when they are used judiciously — a non-trivial task which often results in careless designs that cause screen readers to simply read out long sequences of numbers without any other identifiable information (Fossheim 2020) — they present a fairly low expressive ceiling. The current ARIA specification does not afford rich and nuanced information-seeking opportunities equivalent to those available to sighted readers.

There has been some promising progress for improving support for accessibility within visualization toolkits, and vice-versa for improving native support for charts in screen reader technologies. For instance, Vega-Lite (Satyanarayan et al. 2017) and Highcharts (2021) are beginning to provide ARIA support out-of-the-box. Apple's VoiceOver Data Comprehension feature (Davert and Editorial Team 2019) affords more granular screen reader navigation within the chart, beyond textual summaries and data tables, via four categories of selectable interactions for charts appearing in Apple's Stocks or Health apps. These interactions include Describe Chart, which describes properties of the chart's construction, such as its encodings, axis labels, and ranges; Summarize Numerical Data, which reports min and max data values, and summary statistics like mean and standard deviation; Describe Data Series, which reports the rate-of-change/growth of a curve, trends, and outliers; and Play Audiograph, which plays a tonal representation of the graph's ascending/descending trend over time (Davert and Editorial Team 2019). While Apple's features are presently limited to single-line charts, SAS' Graphics Accelerator (SAS Graphics Accelerator 2018) supports a similar featureset (including sonification, textual descriptions, and data tables) but for a broader range of statistical charts including bar charts, box plots, contour plots, and scatter plot matrices. Our work follows in the spirit of these tools but focuses on web-based visualizations rather than standalone- or platform-integrated software. We go beyond what ARIA supports today to enable high-level and fine-grained screen reader interactions, and hope that our work will help inform ongoing discussions on improving web accessibility standards (e.g., via an Accessibility Object Model (Boxhall et al. 2022).

3 Design Dimensions for Rich Screen Reader Experiences

Currently, the most common ways of making a visualization accessible to screen readers include adding a single high-level textual description (via alt text), providing access to low-level data via a table, or tagging visualization elements with ARIA labels to allow screen readers to step through them linearly (e.g., as with Highcharts 2021). While promising, these approaches do not afford rich information-seeking behaviors akin to what sighted readers enjoy with interactive visualizations. To support systematic thinking about accessible visualization design, we introduce three design dimensions that support rich, accessible reading experiences: *structure*, or how elements of the visualization should be organized for a screen reader to traverse; *navigation*, or the mechanisms by which a screen reader user can move from one element to another; and *description*, or what semantic content the screen reader conveys.

Methods. We began by studying the development of multi-sensory graphical systems, covering work in critical cartography (Wiedel and Groves 1969, Koch 2012), blind education (Aldrich and Sheppard 2001, Godfrey and Loots 2015), tactile graphics (Fujiyoshi et al. 2018, Hasty 2011, de Greef, Moritz, and Bennett 2021, Amick and Corcoran 1997, Butler et al. 2021), and multi-sensory visualization (Hasper et al. 2015, Chundury et al. 2021, Brock et al. 2010, Baker et al. 2016). Drawing on conventions and literature on crip, reflective, and participatory design (Hamraie 2013, Sengers et al. 2005, Costanza-Chock 2020), all authors began an iterative co-design process with Hajas, who is a blind researcher with relevant expertise. Hajas is a screen reader user with a PhD in HCI and accessible science communication, but he is not an expert in visualization research. Co-design — particularly as encapsulated in the disability activism slogan, “*Nothing about us, without us*” (Costanza-Chock 2020) — is important because it can eliminate prototypes that replicate existing tools, solve imaginary problems (i.e., by creating disability dongles (Jackson 2019) or unintentionally produce harmful technology (Shew 2020). To balance engaging disabled users while acknowledging academia’s traditionally extractive relationship with marginalized populations (Cornwall and Jewkes 1995), we intentionally acknowledge Hajas as both co-designer and co-author. We believe that the distinction between co-designer — a phrase that often discounts lived experience as insufficiently academic — and researcher is minimal; technical, qualitative, and experiential expertise are all important components of this research. Hajas’ profile is a perfect example of the intersection between lived experience of existing challenges and solutions, academic experience of research procedures, and an interest in the science of visualization. While he does not represent all screen reader users, his academic expertise and lived experience uniquely qualify him to be both researcher and co-designer. Nevertheless, to incorporate a diverse range of perspectives, we recruited additional participants as part of an evaluative study (§5).

Our work unfolded over 6 months and yielded 15 prototypes. All authors met weekly for hour-long video conferences. In each session, we would discuss the structure and affordances of the prototypes, often by observing and recording Hajas' screen as he worked through them. We would also use these meetings to reflect on how the prototypes have evolved, compare their similarities and differences, and whiteboard potential design dimensions to capture these insights. Following these meetings, Hajas wrote memos detailing the motivations for each prototype, tagging its most salient features, summarizing the types of interactions that were available, enumerating questions that the prototype raises, and finally providing high-level feedback about its usefulness and usability. In the following section, we liberally quote these memos to provide evidence and additional context for our design dimensions.

3.1 Structure

We define *structure* to mean an underlying representation of a visualization that organizes its data and visual elements into a format that can be traversed by a screen reader. Through our co-design process, we identified two components important to analyzing accessible structures: their *form*, or the shape they organize information into; and *entities*, or which parts of the visualization specification are used to translate a chart into a non-visual structure. Design decisions about form and entities are guided by considerations of *information granularity*, or how many levels comprise the range between a high-level overview and individual data values.

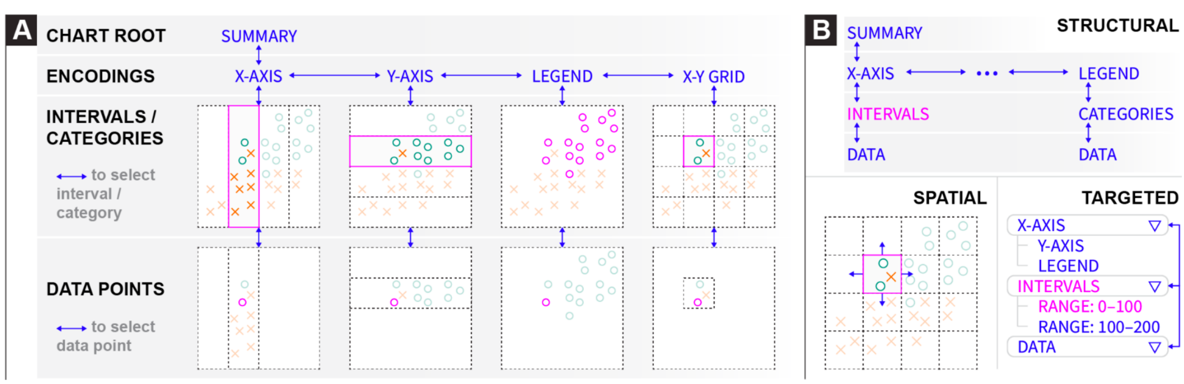

Figure 1: (a) An accessible visualization structure in the form of a tree and comprised of encoding entities. Solid magenta outlines indicate the location of the screen reader cursor. Solid blue arrows between labels indicate available next steps via keyboard navigability (up, down, left, right). (b) Three ways of navigating accessible visualization structures: structural, spatial, and targeted.

Form. Accessible structures organize information about the visualization into different forms, including lists, tables, and trees. Consider existing best practices and common approaches. A rasterized chart with alt text is represented to a screen reader as a single node. SVG-based visualizations can additionally be tagged with ARIA labels to describe the axes, legends, and individual data points. Despite SVG's nesting, screen readers linearize these ARIA labels into a list structure so that the user can step through them sequentially. Data tables, on the other hand, provide a grid structure for screen readers to traverse. At each cell of the grid, the screen reader reads out a different textual description, allowing the user to explore a space by traversing the grid spatially (up, down, left, and right) instead of merely linearly. Accessible visualization research has begun to explore the use of tree structures for storing chart metadata (Weninger et al. 2015), but they remain relatively rare in practice. Trees are our focus here as their branching and hierarchical organization allows users to browse different components of a visualization and traverse them at different levels of detail.

Entities. Where form refers to how nodes in a structure are arranged, entities instead refers to what aspects of the visualization the nodes represent. These aspects can include:

- *Data,* where nodes in the structure represent individual data values or different slices of the data cube (e.g., by field, bins, categories, or interval ranges). For example, in a data table, every node (i.e. cell) represents a data value designated by the row and column coordinates. Depending on the form, data entities can be presented at different levels of detail. For example, one prototype we explored represents a line chart as a binary tree structure (Fig. 2e): the root node represents the entire x-axis domain, and each left and right child node recursively splits the domain in half. Users can traverse the tree downward to binary search for specific values or understand the data distribution.

- *Encodings,* where nodes in the structure correspond to visual channels (e.g., position, color, size) that data fields map to. For instance, consider Figure 1a which depicts the encoding structure of a Vega-Lite scatterplot. The visualization is specified as mappings from data fields to three visual encoding channels: `x`, `y`, and `color`. Thus, the encoding structure, which here takes the form of a tree, comprises a root node that represents the entire visualization and then branches for each encoding channel as well as the data rectangle (x-y grid). Descending into these branches yields nodes that select different categories or interval regions, determined by the visual affordances of the channel. For instance, descending into axis branches yields nodes for each interval between major ticks; x-y grid nodes represent cells in the data rectangle as determined by intersections of the axes gridlines; and legend nodes reflect the categories or intervals of the encoding channel (i.e., for nominal or quantitative data respectively). Finally, the leaves of these branches represent individual data values that fall within the selected interval or category.

- *Annotations,* where nodes in the structure represent the rhetorical devices a visualization author may use to shape a visual narrative or guide reader interpretation of data (e.g., by drawing attention to specific data points or visual regions). Surfacing annotations in the visualization structure allows screen reader users to also benefit from and be guided by the author's narrative intent. For example, Figure 2d illustrates an annotation tree structure derived from an example line chart with two annotations highlighting intervals in the temporal x-axis. The root of the tree has two children representing the two annotated regions. These two annotation nodes have a child node for each data point that is highlighted within the region of interest.

Considerations: Information Granularity. When might users prefer nested structures (i.e. trees) over flat structures (i.e., lists and tables)? Like sighted users, screen reader users seek information by looking for an overview before identifying subsets to view in more detail (Sharif et al. 2021). Trees allow users to read summary information at the top of the structure, and traverse deeper into branches to acquire details-on-demand. Kim et al. (2021) use the term *information granularity* to refer to the different levels of detail at which an accessible visualization might reveal information. They organize granularity into three levels: existence, overview, and detail. *Existence* includes information that a chart is present, but no information about underlying data. *Overview* includes summary information about data — e.g. axes, legends, and summary statistics like min, max, or mean — but not individual data points. *Detail* includes information about precise data values.

We use the root node to signal the existence of the tree, and deeper nodes in the tree reflect finer levels of granularity. Branch nodes give an overview summary about the data underneath, providing information scent (Pirolli and Card 1999), while leaf nodes map to individual data points. In his feedback about the prototype shown in Figure 1, Hajas wrote *"considering how difficult reading a scatterplot with a screen reader is due to its sequential reading nature, the tree structure makes the huge number of data points fairly readable."*

Entities are not mutually exclusive, and a structure might opt to surface different entities in parallel branches. We prototyped a version of Figure 2d which placed an encoding tree and annotation tree as sibling branches under the root node. Users could descend down a given branch, and switch to the equivalent location in the other branch at will. These design decisions are motivated by findings in prior work: by placing encodings and annotations as co-equal branches, we produce a structure that preserves the agency of screen reader users either to start with the narrative arc of annotations, or follow it after having the chance to interpret the data for themselves (Lundgard and Satyanarayan 2021). As Hajas confirms *"Depending on my task, either the encoding or annotation tree could be more important. If my task involved checking population growth in the last 100 years, I would start with the encodings. If I were to look for sudden changes in population numbers, such war-time mortality effects, I would start exploring the annotations, then tunnel back to the other tree."*

3.2 Navigation

Screen reader users need ways to traverse accessible structures to explore data or locate specific points. When browsing a webpage, screen readers provide a cursor that represents the current location in the page. Users use keyboard commands to step the cursor backward and forward in a sequential list of selectable items on the page, or jump to important locations such as headers and links. Through our prototyping process, we developed three ways of navigating through an accessible structure: *structural navigation*, *spatial navigation*, and *targeted navigation* (Fig. 1b). A key concern across these navigation schemes is reducing a user's cognitive load by affording a sense of the boundaries of the structure.

Structural Navigation. Structural navigation refers to ways users move within the accessible structure. We identify two types of structural navigation. *Local navigation* refers to step-by-step movements between adjacent nodes in the structure. This includes moving up and down levels of a hierarchy, or moving side to side between sibling elements. *Lateral navigation* refers to movement between equivalent nodes in adjacent sub-structures. For example, Fig. 2a depicts a multi-view visualization with six facets. When the cursor is on a Y-axis interval for the first facet, directly moving to the same Y-axis interval on the second facet is a lateral move.

Spatial Navigation. Sometimes users want to traverse the visualization according to directions in the screen coordinate system. We refer to this as spatial navigation. For example, when traversing part of an encoding structure that represents the visualization's X-Y grid, a downward structural navigation would go down a level into the currently selected cell of the grid, showing the data points inside the cell. A downward spatial navigation, in contrast, would move to the grid cell below the current one — i.e. towards the bottom of the Y-axis. Spatial navigation is also useful when navigating lists of data points, which may not be sorted by X or Y value in the encoding structure. Where a leftward structural navigation would move to the previous data point in the structure, a leftward spatial navigation would move to the point with the next lowest X value.

Targeted Navigation. Navigating structurally and spatially requires a user to maintain a mental map of where their cursor is relative to where they want to go. If the user has a specific target location in mind, maintaining this mental map in order to find the correct path in the structure to their target can create unnecessary cognitive load. We use targeted navigation to refer to methods that only require the user to specify a target location, without needing to specify a path to get there. For example, the user might open a list of locations in the structure and select one to jump directly there. Screen readers including JAWS and VoiceOver implement an analogous form of navigation within webpages. Instead of manually stepping through the page to find a specific piece of content, users can open a menu with a list of locations in the page. These locations are defined in HTML using ARIA landmark roles, which can designate parts of the DOM as distinct sections when read by a screen reader. When a screen reader user open the list of landmarks and selects a landmark, their cursor moves directly to that element.

Considerations: Boundaries & Cognitive Load. Screen reader users only read part of the visualization at a time, akin to a sighted user reading a map through a small tube [hasty_guidelines_2011]. How do they keep track of where they are? In our co-design process, we found it easiest for a user to remember their location relative to a known starting point, which is corroborated by literature on developing spatial awareness for blind people [wiedel_tactual_1969, li_editing_2019, chundury_towards_2021]. Hajas noted the prevalence of the Home and End shortcuts across applications for returning to a known position in a bounded space (e.g. the start/end of a line in a text editor). We also found that grouping data by category or interval was helpful for maintaining position. Hajas noted that exploring data within a bounded region was like entering a room in a house. In his analogy, a house with many smaller rooms with doors is better than a house with one big room and no doors. Bounded spaces alleviate cognitive load by allowing a user to maintain their position relative to entry points.

Comparing navigation techniques, Hajas noted that spatial felt *"shallow but broad"* while targeted felt *"deep but narrow."* While he expressed a personal preference for deep-narrow structures, he nevertheless *"would not give up [spatial navigation] because it makes me believe I'm actually interacting with a visualization."* This insight demonstrates the value of offering multiple complementary navigation techniques. Moreover, while targeted navigation facilitates quick searching and doesn't require the user to maintain a mental map to find specific data points, structural and spatial exploration enable more open-ended data exploration. It also provides a mechanism for establishing common ground with sighted readers (e.g., allowing both blind and sighted readers to understand a line segment as being "above" or "higher" than another).

3.3 Description

When a user navigates to a node in a structure, the screen reader narrates a description associated with that node. For example, when navigating to the chart's legend, the screen reader output might articulate visual properties of the chart's encoding: *"Category P has color encoding red, Q has color encoding blue"* (Figure 1). Or, if that visual semantic content isn't relevant to understanding the data, it might ignore the color: *"each datum belongs to either Category P or Q."* The content, composition, and verbosity of the description can affect a user's comprehension of the data. Designers must consider context & customization when describing charts.

Content. Semantic *content* is the meaningful information conveyed not only through natural language utterances, but also through the visualization (a graphical language, Bertin 1983). Because graphics convey myriad different kinds of content, the challenge of natural language description is to convey information that is not only commensurate with what the chart expresses via graphical language, but also useful to its readers. Accessible chart description guidelines from WGBH (Gould, O'Connell, and Freed 2008), W3C (2019), and others (Jung et al. 2021) offer prescriptions for conveying specific content for blind readers (such as the chart's title, axis encodings, and noteworthy trends). Lundgard and Satyanarayan (2021) expand the scope of these guidelines with a more general conceptual model of four levels of semantic content: chart construction properties (e.g., axes, encodings, marks, title); statistical concepts and relations (e.g., outliers, correlations, descriptive statistics); perceptual and cognitive phenomena (e.g., complex trends, patterns); and domain-specific insights (e.g., socio-political context relevant to the data).

Decoupling a chart's semantic content from its visual representation helps us better understand what data representations afford for different readers. For instance, Lundgard and Satyanarayan find that what blind readers report as most useful in a chart description is not a straightforward translation of the visual data representation. Specifically, simply listing the chart's encodings is much less useful to blind readers than conveying summary statistics and overall trends in the data (Lundgard and Satyanarayan 2021). As Hajas noted, *"I want to see the global trend, which is why sighted people rely on visualization."* For instance, for a stock market chart the reader *"might see the overview from first to last data points, and then zoom into an outlier in the middle."* These findings suggest opportunities interleave different kinds of content at different levels of a hierarchical structure to yield richer, more useful screen reader navigation. For example, injecting summary statistics (say, the existence of outliers within a particular subcategory of the data) higher up in the chart's tree structure (e.g., at the `legend` encoding node) might afford "scent" for "information foraging" (Pirolli and Card 1999), or further exploration down a particular branch (data subcategory) of the tree. Or, if navigating in a targeted fashion, the user might be afforded the option to directly navigate to outliers without traversing the tree.

Composition. The usefulness of a description depends not only on the content conveyed by its constituent sentences, but also on its *composition*: how those sentences are ordered in relation to each other. For example, during our co-design process, Hajas found that when navigating a chart's tree structure, the screen reader output could quickly become redundant, affecting how quickly and efficiently he could pick out the meaningful information at each node. For instance, the utterance *"Category: P, Point 3 of 15, X = 5, Y = 12"* and the utterance *"x = 5, y = 12, Category: P, Point 3 of 10"* afford significantly different experiences for a user who wishes to quickly scan through individual data points. In the first utterance, the reader immediately receives content that helps to situate them in a broader data context, namely data labeled as Category: A at the `legend` node. In the second utterance, the reader immediately receives datum-specific content that helps to rapidly explore the fine-grained details within that data context. Whether a reader prefers one compositional ordering to another will depend on the task they are attempting to accomplish. As Hajas noted, *"I like the label at the beginning of the information, saying at which level of the tree I am at. It is important for knowing where I am. It is also great that this information is only spoken out when I change level, but not when I navigate laterally."* These compositional choices are highly consequential for readers' experience, when they must repeatedly read nearly-identical utterances while navigating a structure.

Verbosity. Whereas composition refers to the ordering of content, *verbosity* refers to how much content the screen reader conveys. More content is not always better. As Hajas noted of Apple's Data Comprehension feature (Davert and Editorial Team 2019): *"It can sometimes be too much information all at once, if it starts reading out all of the data. This is very difficult if you're interested in some data points that are in the middle. It is very play-or-stop."* Depending on the screen reader software, a user may be afforded control over how much content is conveyed. For instance, JAWS offers high, medium, and low verbosity levels (Freedom Scientific 2021). At higher verbosity the screen reader announces more structural, wayfinding content (e.g. the start and end of regions). For data tables, verbosity configurations can affect whether the table size is read as part of the description, and whether row and column labels are repeated for every cell. Descriptions of nodes in an encoding structure might analogously include information about the path from the root — for example, by reminding the user that they are reading Y-axis intervals. These repetitions can help users remember their location within a structure, but additional verbosity is less efficient for comprehending the data quickly.

Considerations: Context & Customization. Apart from its constituent parts (content, composition, verbosity), a description's usefulness also depends on the *context* in which it is read: namely, the reader's task or intent, and familiarity with the data interface. The same description might be useful in some situations, but relatively useless in others. A reader's information needs are fundamentally context-sensitive. For example, as Hajas noted, when reading a news article, it may be satisfactory to accept a journalist's description of the data on good faith. But, when reviewing scientific research, *"I don't necessarily want to just believe what is said in the text, I want to check and double-check the authors' claims. Go down to the smallest numbers in the analysis. I want to be able to look at the confusion matrix and see if they made a mistake or not."* This targeted verification requires a description to afford users with precise look-up capabilities, in contrast to descriptions that may be generated when browsing or exploring the data.

This context-sensitivity reveals an important aspect of usability: a user's familiarity (or lack thereof) with the data interface. Wayfinding content (e.g., *"Legend. Category: A."*) can help a user remember their location in a structure, and may be useful while they assemble a mental map of the visualization. But, as they become accustomed to the interface and visualization, such descriptions may prove cumbersome. Because user needs depend on their task, preferences, and familiarity, interfaces might afford personalization and customization to facilitate context-sensitive description.

4 Example Gallery

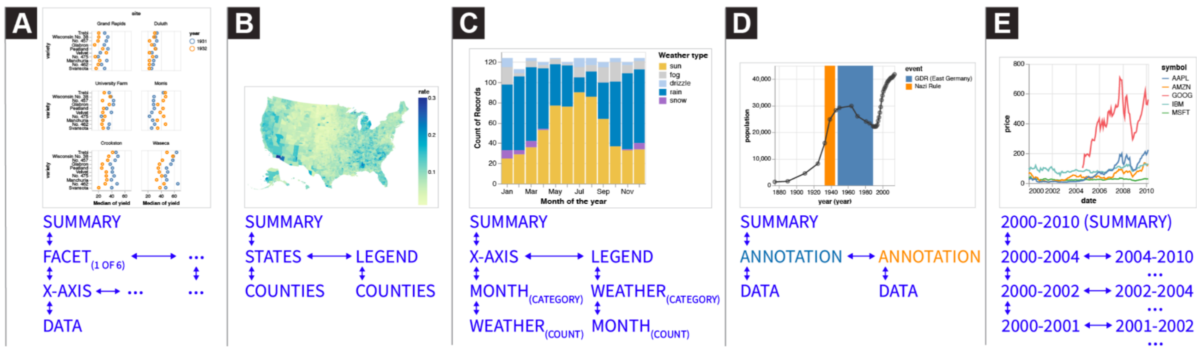

Figure 2: Example structural and navigational schemes generated as part of our co-design process, and applied to diverse chart types.

Our co-design process yielded prototypes that demonstrate a breadth of ways to operationalize our design dimensions. Figure 2 excerpts some of our highest-fidelity prototypes, implemented on top of Vega-Lite (Satyanarayan et al. 2017). As deeply nested structures and dynamic content are not well-supported by ARIA, we implemented our designs as in-memory data structures. Event listeners update the user's position in the structure on keypress, and write text descriptions to an ARIA-Live region (an ARIA role typically used for temporary notifications). To establish common ground with sighted users, we also render the visualization graphically. The user's position in the tree drives a Vega-Lite selection that highlights points when the screen reader user is attending to them.

For every prototype, the `up`, `down`, `left`, and `right` arrow keys enable structural navigation (moving up or down a level, or stepping through siblings respectively). For example, within the facet level of Fig. 2(a), the user can press `left` or `right` keys to move between the six subplots of the multiview chart. On charts that contain a node representing the x-y grid, users can also use the WASD keys to spatially navigate the grid and data points within that branch (mimicking an interaction found in video games).

These prototypes highlight different compositions of structures and navigation schemes. Fig. 2(a) includes `shift+left` and `shift+right` for lateral navigation across facets: pressing these keys at any node within a facet branch will navigate to the same location under an adjacent branch (subplot). With the choropleth (Fig. 2(b)), we group data in the encoding structure by U.S. state; users can then drill down into counties across either this branch or the legend one. Fig. 2(c) offers two different paths for drilling down: month first, or weather first. Fig. 2(d) structures the tree by annotations rather than encoding: users can descend into the time intervals designated by the orange and blue rectangles, and view points within those intervals. Finally, Fig. 2(e) organizes its tree in terms of data, offering a binary search structure through the years.

5 Evaluation

To evaluate our contribution, we conducted 90-minute Zoom studies with 13 blind and visually impaired participants. Participants were asked to explore three prototype accessible screen reader experiences, shown one after another each with a different dataset. The goal of our evaluation was not to determine which particular combination of design elements was "best," but rather to be exploratory — to compare the relative strengths and advantages of instantiations of our design dimensions, and understand how they afford different modes of accessible data exploration.

5.1 Evaluation Setup & Design

Following Frøkjær and Hornbæk's (2005) Cooperative Usability Testing (CUT) method, Zong and Lee conducted each session by alternating between the role of guide (i.e., talking to the user and explaining the prototype) and logger (i.e., keeping track of potential usability problems, interpreting the data to prepare for becoming the guide). We began each session with a semi-structured interview to understand participants' current experiences with data and the methods they use to make inaccessible forms of data representation usable (script included in supplementary material). The rest of the session focused on each of the three prototypes in turn, with each condition split into two phases: interaction and interpretation. In the interaction phase, Zong or Lee guided participants through the prototypes and asked participants to use them and comment on their process, in the style of Hutchinson et al.'s (2003) technology probes. Then, the authors switched roles and began a cooperative interpretation phase, where the authors and participants engaged in a constructive dialogue to jointly interpret the usability problems and brainstorm possible alternatives to the current prototype. In this method, participants influence data interpretation, allowing for more rapid analysis than traditional think-aloud studies as some analysis is built into each evaluation session with instant feedback or explanation from participants (Frøkjær and Hornbæk 2005).

Prototypes. The in-depth nature of our cooperative interpretation sessions required us to balance the total number of prototypes evaluated (so that participants would have time to thoroughly learn and interact with each one) with a time duration appropriate for a Zoom session (limited to 90 minutes to avoid exhausting participants). Accordingly, we selected the following three prototypes, each representing a different aspect of our design dimensions:

- **Table:** An accessible HTML data table with all rows and three columns from the classic Cars dataset, in order to compare our prototypes with existing accessibility best practice.

- **Multi-view:** Becker's barley yield trellis display (Becker, Cleveland, and Shyu 1996) as shown in Fig. 2a. This prototype features local and lateral structural navigation via the arrow keys and with the shift modifier respectively, as well as spatial navigation via WASD.

- **Target:** A single-view scatterplot, illustrated in Fig. 1, depicting the Palmer Penguins dataset (Horst, Hill, and Gorman 2020). In addition to structural and spatial navigation, targeted navigation is available via three dropdown menus corresponding to the structural levels.

The table is our control condition, as it follows existing best practice for making data accessible to screen readers. multi-view enables us to study how users move between levels of detail, and whether they could navigate and compare small multiple charts. Finally, target allows us to compare how and when our participants use the three different styles of navigation (structural, spatial, and targeted). We presented the prototypes in this sequence to all participants to introduce new features incrementally.

Participants. We recruited 13 blind and visually impaired participants through our collaborators in the blind community and through a public call on Twitter. Each participant received $50 for a 90-minute Zoom session. We provide aggregate participant data following ethnographic practice to protect privacy and not reduce participants to their demographics (Saunders, Kitzinger, and Kitzinger 2015). Half of our participants were totally blind (n=7), while others were almost totally blind with some light perception (n=4) or low vision (n=2). Half of them have been blind since birth (n=7). Participants were split evenly between Windows/Chrome (n=7) and Mac/Safari (n=6). Windows users were also split evenly between the two major screen readers (JAWS, n=3; NVDA, n=4), while all Mac participants used Apple VoiceOver. These figures are consistent with recent surveys conducted by WebAIM which indicate that JAWS, NVDA, and VoiceOver are the three most commonly used screen readers [noauthor_screen_2021]. Demographically, 70% of our participants use he/him pronouns (n=9) and the rest use she/her pronouns (n=4). One participant was based in the UK while the rest were spread across eight US states. Participants self-reported their ethnicities (Caucasian/white, Asian, and Black/African, Hispanic/Latinx), represented a diverse range of ages (20–50+) and had a variety of educational backgrounds (high school diploma through to undergraduate, graduate, and trade school degrees). Nine participants reported themselves as slightly or moderately familiar with statistical concepts and data visualization methods, 2 as expertly familiar, and one as not at all familiar. Five participants described data analysis and visualization tools as an important component in their professional workflows, and 8 interacted with data or visualizations more than 1–2 times/week.

5.2 Quantitative Results

PROMPT : *When using this prototype ...*	TASK [BM13]	TABLE	MULTI-VIEW	TARGETED
How enjoyable was it to interact with the data?	`enjoy`	**3** [3.31] (0.95)	**4** [3.77] (1.01)	**4** [3.54] (0.97)
How easy was it to generate and answer questions?	`discover`	**4** [3.15] (1.34)	**3** [3.00] (1.08)	**3** [3.23] (1.17)
If you already knew what information you were trying to find, how easy would it be to look up or locate those data?	`lookup-locate`	**3** [3.31] (1.32)	**4** [3.77] (1.17)	**4** [3.38] (1.19)
If you didn't already know which information you were trying to find, how easy would it be to browse or explore the data?	`browse-explore`	**2** [3.00] (1.68)	**2** [2.69] (1.11)	**3** [3.00] (1.29)
PROMPT : *When using this prototype ...*	USE [KS99]	TABLE	MULTI-VIEW	TARGETED
How easy was it to learn to use?	ease-of-use	**4** [4.15] (0.99)	**3** [2.69] (0.75)	**3** [3.15] (1.34)
How useful would it be to have access to this interaction style for engaging with data?	perceived usefulness	**4** [4.15] (0.80)	**4** [4.00] (0.82)	**4** [4.15] (1.07)

Table 1: Rating scores for each prototype (Table, Multi-view, Targeted) on a five point Likert scale where 1 = Very Difficult (Very Unenjoyable) and 5 = Very Easy (Very Enjoyable). Median scores are shown in boldface, averages in brackets, standard deviations in parentheses.

To supplement the cooperative interpretation sessions, participants rated each prototype using a series of Likert questions. We designed a questionnaire with six prompts measuring a subset of Brehmer and Munzner's multi-level typology of abstract visualization tasks (Brehmer and Munzner 2013). This framework, however, required some adaptation for non-visual modes of search. In particular, searching with a screen reader requires a sequential approach to data that is at odds with the "at-a-glance" approach sighted readers take to browsing and exploring data. As our prototypes focus on navigation through charts, we collapsed the location dimension of Brehmer and Munzner's search decomposition resulting in two prompts that jointly measure `lookup-locate` and `browse-explore`. We formulated additional questions to measure Brehmer and Munzner's `discover` and `enjoy` tasks as well as more traditional aspects of technology acceptance including *ease-of-use* and *perceived usefulness* (Karahanna and Straub 1999). Participants responded on a five point scale where 1 = Very Difficult/Unenjoyable and 5 = Very Easy/Enjoyable.

Table 1 displays the questionnaire prompts, their corresponding tasks, and statistics summarizing the participants' ratings. A Friedman test found a significant rating difference for the ease-of-use of the prototypes $\chi^2(2, N = 13) = 15.05$, $p < 0.01$, with a large effect size (Kendall's $W = 0.58$). Follow-up Nemenyi tests revealed that *multi-view* was more difficult to use than *table* with statistical significance ($p < 0.01$), but target was not. Additional tests for the other prompts found neither statistically significant differences, nor large effect sizes, between the prototypes. However, median scores (which are more robust to outliers than means, Morris et al. 2018) suggest that participants generally `enjoy` interacting with multi-view and target more, and found them easier to `lookup` or `locate` data with. Moreover, *target* had the highest median score for affording `browse` or `explore` capabilities. Conversely, the *table* was easiest to learn to use, and generally made it easy to `discover`, or ask and answer questions about the data. Notably, in response to the question *"How useful would it be to have access to this interaction style for engaging with data?",* participants on average ranked all prototypes as more-than-useful (*med* = 4, $\mu \geq 4$). These statistics provide only a partial picture of participants' experiences with the prototypes (Bagozzi 2007). Thus, we elucidate and contextualize reasons behind their scores through qualitative analysis.

5.3 Qualitative Results

After the interviews, we qualitatively coded the notes taken by the logger with a grounded theory approach (Charmaz 2006). We performed open coding in parallel with the interviews (i.e., coding Monday's interviews after finishing Tuesday's interviews). We then synthesized the codes into memos, from which we derived these themes.

Tables are familiar, tedious, but necessary. Every participant noted that tables were their primary way of accessing data and visualizations. While tables are an important accessible option, participants overwhelmingly reported the same problems: they are ill-suited for processing large amounts of data and impose high cognitive load as users must remember previous lines of the table in order to contextualize subsequent values. As P2 reported, *"if I'm trying to get a general sense of the table, I'll just scroll through and see what values there are. But there's 393 rows, so I'll never scroll through all of it…I can't really get a snapshot."* P11 said that *"Finding relationships can be tricky if you're in a table, because you've got to either have a really good memory or just get really lucky. […] If I didn't know what I was looking for, forget it."* At most, participants tabbed through 20–30 rows during our sessions, but did so only because of the questions we posed (e.g., "is there a relationship between horsepower and mileage?") and noted that if they encountered this table outside of the study, they would tab past a few rows to check for summary statistics and then move on.

While it is not enjoyable to explore or build a mental model of data with static tables, participants still emphasized their necessity because of the format's familiarity: *"in terms of accessibility, tables are infinitely more useful because there is a standard way of navigating them in whatever your preferred screen reader is. With different representations, a blind person may not be trained to interpret it"* (P2). This builds on prior literature (Sharif et al. 2021) and echoes testimony from participants who had some difficulty with the new prototypes; they reported that they lacked expertise and therefore found it difficult to work with non-tabular data (P8, 10). In other words, to maximize accessibility, it is crucial to include a table view of the data *in addition* to other forms of novel interaction.

Prior exposure to data analysis and representations increases the efficacy of spatial representations. Participants who had experience conducting data analysis or reading tactile graphs/maps were able to easily develop a spatial understanding of how each prototype worked. Five participants (P2–4, 11, 13) made direct connections between the multi-view and target prototypes, and the tactile graphs they encountered in school. Three participants (P2, 11, 12) found their software engineering experience made it easier to understand and navigate the prototypes' hierarchical structure. Previous literature on tactile mapping has also shown how developing tactile graphical literacy is crucial for building spatial knowledge, but they emphasize that it is not sufficient for being able to conduct and understand data analysis (Hasty 2011, Godfrey and Loots 2015). Since our participants already had an existing spatial framework, it became easier to explain how a prototype might work using their prior experience as a benchmark, which has been corroborated by similar studies in tactile cartography (Weidel and Groves 1969, Aldrich and Sheppard 2001, Sheppard and Aldrich 2001). Importantly, our participants were able to find specific origin points that they could return to in order to navigate the different branches of the tree, which would be further aided with help menus and mini-tutorials to understand the keyboard shortcuts (P2). Being able to shift between origin points is especially important for switching between graphs or between variables. By contrast, participants who had more difficulty with the prototypes (P8, 10) pointed to their lack of experience working with non-tabular data. P10 reported that being able to mentally visualize data points within a grid was a specific challenge. *"I suspect that this might be understandable to someone who's done this before,"* he said, *"I don't do well with these charts unless they're converted back into tables."*

Structure: **Hierarchical representations make it possible to effectively convey insights with minimal cognitive load.** While static tables are the most common accessible option to interactive visualizations, eight of our participants (P2–5, 7, 10, 11, 13) expressed a desire to filter and sort the data so that they could begin to explore possible trends without wading line by line. Sorting and filtering a table is one way to look for trends but, to get a summary view of the data quickly, a system must provide snapshots in smaller intervals so that users can easily construct a larger picture or choose specific slices of the data to explore further (i.e., "details on demand"; Shneiderman 2003; Kim et al. 2021). With multi-view and target P4 said, *"I always want more layers and details, but some charts had too much…This was a happy medium between having the information I wanted and presenting it in a way that I can keep up with."* P5 also noted that he liked *"having the ability to scroll through at a higher level and then drill down deeper if that's of interest."* By giving users a way to quickly skip through the data across specific axes, they are able to rapidly generate a broader mental image of each graph and drill down further to collect more details. *"When I was working with the table, I [started building] a table in my head,"* P2 shared. *"I had a rough representation of it as a scatter plot. But here, I know how to drill down and up between different layers of data grids, so that I can get the overall picture… [It gives me] different ways of thinking."* Being able to control the parts of the data that were most important to them was also an issue of trust, as it also provided a way for users to reach conclusions for themselves rather than rely on the interpretation of others: *"It's hard to mix…doing your own analysis and be given a text description that you have to just trust"* (P12). In their own workflows, these participants reported downloading static tables to further examine and manipulate with Excel, which they would use to create summary statistics or intervals to move more quickly through the data.

Navigation: **Reading a visualization with a screen reader entails constant hypothesis testing and pattern-making.** Since screen reader users parse data iteratively, nine of our participants (P1–5, 7, 8, 11, 13) described reading a visualization as a process of slowly building up a mental model and constantly testing it to see where the patterns may no longer hold. "I'm going row by row, not memorizing exact numbers but building a pattern in my head, and looking at the other rows to test my theory," reported P3. In other words, our participants engaged in a continuous state of updating and validating (Munzner 2009) their mental images as new data challenged the existing patterns they have pieced together. multi-view and target accelerated this process, as participants were able to more rapidly identify specific components that they wanted to test. For example, P2 intentionally moved quickly across each level of the structure hoping to find its "edges," or the minimum and maximum limits of each axis and grid. *"Visually, it might look like I'm doing a lot of jumping around,"* he said, *"[but it's] because I'm trying to build the picture in a way that makes sense for me."* Similarly, P5 started building his mental model of the visualization by drilling up and down the grid to create a spatial image of the data: *"I'm thinking more in spatial terms just because [this] is a new method of navigating to me. […] I'm moving through the grid…I'm thinking of drilling down into that square to get more information."*

Target made it especially easy for participants to test their hypotheses by giving them direct access to components that might break their hypotheses. P5 reported that it allowed him to *"navigate to areas…that I'm interested in, skipping over stuff that's not of interest,"* and P4 likened it to *"[being] able to go directly to what you want in a grocery inventory rather than going through each item one by one."* The ability to use structural, spatial, and target navigation in both multi-view and target respectively facilitated the hypothesis-testing and pattern-making behaviors that our participants were accustomed to with static tables, and gave them an additional mental model for working with the data. As P1 noted, these prototypes gave her a richer understanding of the data by helping her piece together *"both the picture and the mathematical pattern,"* whereas table afforded only the latter.

Description: **Cursors and roadmaps are important for understanding where you are.** Being able to capture both a high-level overview of the information while preserving the ability to drill down into the data is a crucial component to accessing interactive visualizations (Sharif et al. 2021). To navigate between these two levels, however, our participants emphasized the importance of markers to help them understand where they could move. target addressed this with dropdown menus that allowed participants to navigate to any part of the visualization, explore, and then return to where they had started. In the words of P4, *"[This] mode is freedom for the user. Being able to jump around and move in real time as you would with your hand gives you a new way of exploring the information."* Multi-view approached this issue by allowing participants to move throughout the grid. *“With the table, I was trying to hold the numbers in my head and I wasn’t trying to visualize it or anything,”* said P3. *“With [Multi-View], I can sort of think about it more like a visualization since I can move up and down, left and right. Even though I can use the arrows in the table, it just doesn’t feel the same. I’m still feeling around and seeing what I can find.”* Without these navigation tools, P7 noted that *“It’s too easy to get lost …I don’t know how to backtrack.”* To orient herself, P13 would first test to see if she was at the corner cells in the visualizations (e.g., *“Am I in the upper left or the bottom right cell here?”*) so that she could contextualize her position within the visualization and return to a point of origin. *“I know that I must be at the bottom left cell here because I can’t go to the left,”* P13 said, *“but being able to know where that is beforehand would be very helpful.”*

6 Discussion and Future Work

In this paper, we explore how structure, navigation, and description compose together to yield richer screen reader experiences for data visualizations than are possible via alt text, data tables, or the current ARIA specification. Our results suggest promising next steps about accessible interaction and representation for visualizations.

6.1 Enabling Richer Screen Reader Experiences Today

Although our design dimensions highlight a diverse landscape of screen reader experiences for data visualizations, our study participants attested to the value of following existing best practices. Namely, alt text and data tables provide a good baseline for making visualizations accessible. Thus, visualization authors should consider adopting our design dimensions to enable more granular information access patterns only after these initial pieces are in place.

Existing visualization authoring methods, however, are likely insufficient for instantiating our design dimensions or producing usable experiences for screen reader users. In particular, it currently falls entirely on visualization authors to handcraft appropriate structures, navigational techniques, and description schemes on a per-visualization basis. As a result, besides being a time-consuming endeavor, idiosyncratic implementations can introduce friction to the reading process. For instance, per-visualization approaches might not account for an individual user’s preferences in terms of verbosity, speed, or order of narrated output — three properties which varied widely among our study participant in ways that did not correlate with education level or experience with data. Thus, to scale and standardize this process, some responsibility for making visualizations screen reader accessible must be shared by toolkits as well. For example, our prototypes suggest a strategy for translating visualization specifications into hierarchical encoding structures (i.e., encoding channels as individual branches, and using visual affordances such as axis ticks and grid lines to populate the hierarchy levels). If toolkits provide default experiences out-of-the-box, visualization authors can instead focus on customizing them to be more meaningful for their specific visualization, and screen reader users have a stronger guarantee that the resultant experiences will be more usable and respectful of their individual preferences.

Current web accessibility standards also present limitations for realizing our design dimensions. For instance, while ARIA roles would seem a natural place to start, we were unable to use them while controlling navigational boundaries for hierarchical structures (§3.2). Similarly, there is no standard way to determine which element the screen reader cursor is selecting. Where ARIA has thus far focused on annotating documents with the semantics of a pre-defined palette of widgets, future web standards might in- stead express how elements respond to the interaction affordances of screen readers. For example, ARIA could offer explicit support for overview/detail hierarchies and different levels of description detail that can be progressively read according to user preferences.

6.2 Studying and Refining the Design Dimensions

Our conversations with study participants also helped highlight that design considerations can differ substantially for users who are totally blind compared to those who have low-vision. For example, partially-sighted participants used screen magnifiers alongside screen readers. As a result, they preferred verbose written descriptions alongside more terse verbal narration. Magnifier users also wished for in situ tooltips, which would eliminate the need to scroll back and forth between points and axes to understand data values. However, promisingly, we found that using a screen reader and magnifier together affords unique benefits: *"I would have missed this point visually if I solely relied on the magnifier because the point is hidden behind another point"* (P12). Future work should more deeply explore how accommodations might complement and conflict when designing for different kinds of visual disability.

Similarly, in scoping our focus to screen readers and, thus, text-to-speech narration, we refrained from considering multi-sensory modalities in our design dimensions. Yet, we found that most participants had previous experience with multi-sensory visualization, including sonification (P5, 7, 9, 13), tactile statistical charts (P2–4, 10, 11, 13), and haptic graphics (P3, 4, 11, 13). Some participants reported that a combination of modalities would further enhance their experience — for example, getting a sonic overview of a line chart before reading more detailed text descriptions. Other participants, however, cautioned that adding multiple modalities can create additional confusion. For example, P7 noted that *"There's often a lack of explanation about how to map between sound and text."* Based on this testimony, it is unlikely that "sensory modalities" are merely an additional, independent dimension within our framework. Rather, future work must unpack the affordances of individual modalities, how they interact with one another, and how they impact the design of structure, navigation, and description.

6.3 What are Accessible Interactions for Data Visualizations?

In visualization research, we typically distinguish between static and interactive visualizations, where the latter allows readers to actively manipulate visualized elements or the backing data pipeline (Yi et al. 2007, Heer and Shneiderman 2012). Screen readers, however, complicate this model: reading is no longer a process that occurs purely “in the head” but rather becomes an embodied and interactive experience, as screen reader users must intentionally perform actions with their input devices in order to step through the visualization structure. While some aspects of this dichotomy may still hold, it is unclear how to cleanly separate static reading from interactive manipulation in the context of screen reader accessible visualizations, if these notions are conceptually separable at all. For instance, Hajas likened the navigation our prototypes afforded to *“shifting eye gaze, shifting focus of perceptual attention. When I navigate a visualization, naturally I would say ‘I’m looking at this figure’ and not that ‘I’m interacting with this figure’.”* Analogously, recent results in graphical perception find that sighted readers do not simply “see” visualizations in a single glance but rather perform active visual filtering operations (Boger, Most, and Franconeri 2021). However, when using the binary tree prototype (Fig. 2e), Hajas noted a more distinct shift from reading to interacting. He said, *“it gave me the impression that I’m not just looking selectively, but I focus and zoom into the data,”* analogously to zoom interactions that change the viewport for sighted readers. Better characterizing the shift that occurs with this prototype, and exploring accessible manipulations of visualizations that allow screen reader users to meaningfully conduct data analysis, are compelling opportunities for future work.

References for Chapter 4

Ablart, Damien, Carlos Velasco, and Marianna Obrist. "Integrating Mid-Air Haptics into Movie Experiences." In *Proceedings of the 2017 ACM International Conference on Interactive Experiences for TV and Online Video*, 77–84. TVX '17. New York, NY, USA: ACM, 2017. https://doi.org/10.1145/3077548.3077551.

Abreu, Stênio, Jorge Sassaki Resende Silva, Graciela Pires dos Anjos, Hérlon Manollo Cândido Guedes, Soraia Silva Prietch, Paula Christina Figueira Cardoso, and André Pimenta Freire. "Usability Evaluation of a Resource to Read Mathematical Formulae in a Screen Reader for People with Visual Disabilities." In *Proceedings of the 18th Brazilian Symposium on Human Factors in Computing Systems*, 1–11. IHC '19. New York, NY, USA: Association for Computing Machinery, 2019. https://doi.org/10.1145/3357155.3358480.

Abu Doush, Iyad, Enrico Pontelli, Tran Cao Son, Dominic Simon, and Ou Ma. "Multimodal Presentation of Two-Dimensional Charts: An Investigation Using Open Office XML and Microsoft Excel." *ACM Transactions on Accessible Computing* 3, no. 2 (November 1, 2010): 8:1-8:50. https://doi.org/10.1145/1857920.1857925.

"Accessible Data Science Beyond Visual Models: Non-Visual Interactions with R and RStudio Packages," January 21, 2021. https://rstudio.com/resources/rstudioglobal-2021/accessible-data-science-beyond-visual-models-non-visual-interactions-with-r-and-rstudio-packages/.

Ahmetovic, Dragan, Nahyun Kwon, Uran Oh, Cristian Bernareggi, and Sergio Mascetti. "Touch Screen Exploration of Visual Artwork for Blind People." In *Proceedings of the Web Conference 2021*, 2781–91. WWW '21. New York, NY, USA: Association for Computing Machinery, 2021. https://doi.org/10.1145/3442381.3449871.

Albert, Mannheimer. Braille writing device. United States US2850812A, filed July 29, 1955, and issued September 9, 1958. https://patents.google.com/patent/US2850812A/en.

Aldrich, Frances K., and Linda Sheppard. "Tactile Graphics in School Education: Perspectives from Pupils." *British Journal of Visual Impairment* 19, no. 2 (May 1, 2001): 69–73. https://doi.org/10.1177/026461960101900204.

American Association of Workers for the Blind, Association for Education of the Visually Handicapped, and National Braille Association. *The Nemeth Braille Code for Mathematics and Science Notation 1972 Revision*. Louisville, KY: American Printing House for the Blind, 1972.

Amick, Nancy, and Jane Corcoran. "Guidelines for Design of Tactile Graphics." American Printing House for the Blind, 1997. https://www.aph.org/research/guides/.

Baker, Catherine M., Lauren R. Milne, Ryan Drapeau, Jeffrey Scofield, Cynthia L. Bennett, and Richard E. Ladner. "Tactile Graphics with a Voice." *ACM Transactions on Accessible Computing* 8, no. 1 (January 13, 2016): 3:1-3:22. https://doi.org/10.1145/2854005.

Barwich, Ann-Sophie. "Measuring the World: Olfaction as a Process Model of Perception," January 1, 2018.

Bergman, Eric, and Earl Johnson. "Towards Accessible Human-Computer Interaction." In *Advances in Human-Computer Interaction*, 87–113, 1995.

Bettinger, Eliza. "LibGuides: Data Visualization: Methods, Tools, Resources: Accessibility and Design." Accessed May 20, 2021. https://guides.library.cornell.edu/c.php?g=898087&p=6489216.

Billie L. Bentzen, Richard M. Jackson, and Alec F. Peck. "Information about Visual Impairment for Architects and Transit Planners." Vol. II of Final Report. Washington, D.C.: U.S. Department of Transportation, Urban Mass Transportation Administration, 1981. http://archive.org/details/informationabout02bill.

Brock, Anke, Philippe Truillet, Bernard Oriola, and Christophe Jouffrais. "Usage of Multimodal Maps for Blind People: Why and How." In *ACM International Conference on Interactive Tabletops and Surfaces*, 247–48. ITS '10. New York, NY, USA: Association for Computing Machinery, 2010. https://doi.org/10.1145/1936652.1936699.

Buccafusco, Christopher. "Incentivizing Accessible Design." SSRN Scholarly Paper. Rochester, NY: Social Science Research Network, December 3, 2019. https://papers.ssrn.com/abstract=3497902.

Butler, Matthew, Leona M Holloway, Samuel Reinders, Cagatay Goncu, and Kim Marriott. "Technology Developments in Touch-Based Accessible Graphics: A Systematic Review of Research 2010-2020." In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–15. 278. New York, NY, USA: Association for Computing Machinery, 2021. https://doi.org/10.1145/3411764.3445207.

Card, Stuart K, Jock D Mackinlay, and Ben Shneiderman. *Readings in Information Visualization: Using Vision to Think*. San Francisco, Calif.: Morgan Kaufmann Publishers, 1999.

Carvalho, Lucas Pedroso, Raquel Oliveira Prates, and André Pimenta Freire. "A Proposal to Adapt the Semiotic Inspection Method to Analyze Screen Reader Mediated Interaction." In *Proceedings of the 18th Brazilian Symposium on Human Factors in Computing Systems*, 1–11. IHC '19. New York, NY, USA: Association for Computing Machinery, 2019. https://doi.org/10.1145/3357155.3358455.

Choi, Jinho, Sanghun Jung, Deok Gun Park, Jaegul Choo, and Niklas Elmqvist. "Visualizing for the Non-Visual: Enabling the Visually Impaired to Use Visualization." *Computer Graphics Forum* 38, no. 3 (June 2019): 249–60. https://doi.org/10.1111/cgf.13686.

Coppin, Peter. "What Is Lost in Translation from Visual Graphics to Text for Accessibility," 2015, 447–52.

Deque. "Creating Accessible SVGs," January 10, 2019. https://www.deque.com/blog/creating-accessible-svgs/.

"Data Visualizations | Visual Design | Accessibility for Teams." Accessed May 20, 2021. https://accessibility.digital.gov/visual-design/data-visualizations/.

"DOTSYS III: A Portable Program for Braille Translation." Library of Congress, Washington, D.C., Division for the Blind and Physically Handicapped, October 2, 1975.

Elmqvist, Niklas. "Visualizing for the Non-Visual." Medium, April 24, 2019. https://medium.com/hcil-at-umd/visualizing-for-the-non-visual-96a3ee229718.

Erhardt, Robert J., and Michael P. Shuman. "Assistive Technologies for Second-Year Statistics Students Who Are Blind." *Journal of Statistics Education* 23, no. 2 (July 1, 2015): null. https://doi.org/10.1080/10691898.2015.11889733.

Visualising Data. "Five Ways to... Design for Red-Green Colour-Blindness," August 7, 2019. https://www.visualisingdata.com/2019/08/five-ways-to-design-for-red-green-colour-blindness/.

Fossheim, How (not) to make accessible data visualizations, illustrated by the US presidential election by Sarah L. "How (Not) to Make Accessible Data Visualizations, Illustrated by the US Presidential Election. by Sarah L. Fossheim." Accessed May 20, 2021. https://fossheim.io/.

Fossheim, Sarah L. "An Intro to Designing Accessible Data Visualizations." Accessed May 20, 2021. https://fossheim.io/.

———. "An Introduction to Accessible Data Visualizations with D3.Js." Accessed May 20, 2021. https://fossheim.io/.

Fossheim, Tools for designing good-looking accessible interfaces by Sarah L. "Tools for Designing Good-Looking Accessible Interfaces by Sarah L. Fossheim." Accessed May 20, 2021. https://fossheim.io/.

Franconeri, Steven L., Lace M. Padilla, Priti Shah, Jeffrey M. Zacks, and Jessica Hullman. "The Science of Visual Data Communication: What Works." *Psychological Science in the Public Interest* 22, no. 3 (December 2021): 110–61. https://doi.org/10.1177/15291006211051956.

Fujiyoshi, Mamoru, Akio Fujiyoshi, Hiroshi Tanaka, and Toru Ishida. “Universal Design Tactile Graphics Production System BPLOT4 for Blind Teachers and Blind Staffs to Produce Tactile Graphics and Ink Print Graphics of High Quality.” In *Computers Helping People with Special Needs*, edited by Klaus Miesenberger and Georgios Kouroupetroglou, 10897:167–76. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2018. https://doi.org/10.1007/978-3-319-94274-2_23.

Future Date. *Translating Maps and Geographic Data to Non-Visual Formats - Elizabeth Pyatt*. Accessed July 23, 2021. https://www.youtube.com/watch?v=PxJ0Tf9ObW4.

Gibson, Becky. “Accessibility Testing by People with Disabilities.” *24 Accessibility* (blog), December 8, 2019. https://www.24a11y.com/2019/accessibility-testing-by-people-with-disabilities/.

Godfrey, A. Jonathan R., and M. Theodor Loots. “Advice From Blind Teachers on How to Teach Statistics to Blind Students.” *Journal of Statistics Education* 23, no. 3 (November 1, 2015): null. https://doi.org/10.1080/10691898.2015.11889746.

Holmes, Peter. “Making Data Visualization Accessible — a Case Study.” Medium, January 9, 2020. https://uxdesign.cc/making-data-visualization-accessible-a-case-study-e5fb41ac62ad.

Humphreys, Sal. “The Economies within an Online Social Network Market : A Case Study of Ravelry.” In *ANZCA 09 Annual Conference : Communication, Creativity and Global Citizenship*, 2009. http://www.anzca.net/documents/2009-conf-papers/34-the-economies-within-an-online-social-network-market-a-case-study-of-ravelry-1.html.

Hutchins, Edwin L, James D Hollan, and Donald A Norman. “Direct Manipulation Interfaces.” *Human-Computer Interaction* 1 (1985): 311–38.

Hutchinson, Hilary, Benjamin B Bederson, Allison Druin, Catherine Plaisant, Wendy Mackay, Helen Evans, Heiko Hansen, et al. “Technology Probes: Inspiring Design for and with Families,” n.d., 8.

Kearney-Volpe, Claire, and Amy Hurst. “Accessible Web Development: Opportunities to Improve the Education and Practice of Web Development with a Screen Reader.” *ACM Transactions on Accessible Computing* 14, no. 2 (July 21, 2021): 8:1-8:32. https://doi.org/10.1145/3458024.

Keen Blog. “Accessibility Considerations in Data Visualization Design.” *Keen* (blog), July 16, 2019. https://keen.io/blog/accessibility-in-data-vis/.

Keyes, Os, Burren Peil, Rua M. Williams, and Katta Spiel. "Reimagining (Women's) Health: HCI, Gender and Essentialised Embodiment." *ACM Transactions on Computer-Human Interaction* 27, no. 4 (September 25, 2020): 1–42. https://doi.org/10.1145/3404218.

Kim, N. W., S. C. Joyner, A. Riegelhuth, and Y. Kim. "Accessible Visualization: Design Space, Opportunities, and Challenges." *Computer Graphics Forum* 40, no. 3 (June 2021): 173–88. https://doi.org/10.1111/cgf.14298.

Klemmer, Scott R., Björn Hartmann, and Leila Takayama. "How Bodies Matter: Five Themes for Interaction Design." In *Proceedings of the 6th ACM Conference on Designing Interactive Systems - DIS '06*, 140. University Park, PA, USA: ACM Press, 2006. https://doi.org/10.1145/1142405.1142429.

Krufka, Stephen Edward. "Visual to Tactile Conversation Theory." M.S. thesis, University of Delaware, 2006.

Li, Jingyi, Son Kim, Joshua A. Miele, Maneesh Agrawala, and Sean Follmer. "Editing Spatial Layouts through Tactile Templates for People with Visual Impairments." In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, 1–11. Glasgow, Scotland Uk: ACM Press, 2019. https://doi.org/10.1145/3290605.3300436.

MacLeod, Haley, Cynthia L. Bennett, Meredith Ringel Morris, and Edward Cutrell. "Understanding Blind People's Experiences with Computer-Generated Captions of Social Media Images." In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 5988–99. Denver Colorado USA: ACM, 2017. https://doi.org/10.1145/3025453.3025814.

Marson, Stephen M., Charles F. Harrington, and Adam Walls. "Teaching Introductory Statistics to Blind Students." *Teaching Statistics: An International Journal for Teachers* 35, no. 1 (2013): 21–25. https://doi.org/10.1111/j.1467-9639.2012.00510.x.

Martínez, Rubén Alcaraz, Mireia Ribera Turró, and Toni Granollers Saltiveri. "Accessible Statistical Charts for People with Low Vision and Colour Vision Deficiency." In *Proceedings of the XX International Conference on Human Computer Interaction*, 1–2. Interacción '19. New York, NY, USA: Association for Computing Machinery, 2019. https://doi.org/10.1145/3335595.3335618.

Mazanec, Melanie. "Data Visualization Accessibility: Where Are We Now, and What's Next?" Medium, July 6, 2020. https://medium.com/nightingale/data-visualization-accessibility-where-are-we-now-and-whats-next-b2c9eeac4e8b.

McDonnell, Emma, Leah Findlater, and Audrey Desjardins. "Bridging Accessibility and Disability Studies Research to Speculate Beyond Ability Norms," n.d., 6.

Miesenberger, Klaus, Joachim Klaus, and Wolfgang Zagler, eds. *Computers Helping People with Special Needs: 8th International Conference, ICCHP 2002, Linz, Austria, July 15-20, 2002: Proceedings*. Lecture Notes in Computer Science, Lecture Notes in Artificial Intelligence 2398. Berlin ; New York: Springer, 2002.

Obrist, Marianna, Rob Comber, Sriram Subramanian, Betina Piqueras-Fiszman, Carlos Velasco, and Charles Spence. "Temporal, Affective, and Embodied Characteristics of Taste Experiences: A Framework for Design." In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems - CHI '14*, 2853–62. Toronto, Ontario, Canada: ACM Press, 2014. https://doi.org/10.1145/2556288.2557007.

Obrist, Marianna, Sue Ann Seah, and Sriram Subramanian. "Talking about Tactile Experiences." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, 1659. Paris, France: ACM Press, 2013. https://doi.org/10.1145/2470654.2466220.

Obrist, Marianna, Sriram Subramanian, Elia Gatti, Benjamin Long, and Thomas Carter. "Emotions Mediated Through Mid-Air Haptics." In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2053–62. CHI '15. New York, NY, USA: ACM, 2015. https://doi.org/10.1145/2702123.2702361.

Obrist, Marianna, Alexandre Tuch, and Kasper Hornbæk. "Opportunities for Odor: Experiences with Smell and Implications for Technology." *Conference on Human Factors in Computing Systems - Proceedings*, April 26, 2014. https://doi.org/10.1145/2556288.2557008.

Offenhuber, Dietmar. "Data by Proxy -- Material Traces as Autographic Visualizations." *ArXiv:1907.05454 [Cs, Math]*, July 11, 2019. http://arxiv.org/abs/1907.05454.

Paixão Silva, Luiz Felipe da, Otavio de Faria Oliveira, Evelise Roman Corbalan Góis Freire, Rosana Maria Mendes, and André Pimenta Freire. "How Much Effort Is Necessary for Blind Users to Read Web-Based Mathematical Formulae? A Comparison Using Task Models with Different Screen Readers." In *Proceedings of the XVI Brazilian Symposium on Human Factors in Computing Systems*, 1–10. IHC 2017. New York, NY, USA: Association for Computing Machinery, 2017. https://doi.org/10.1145/3160504.3160549.

Paterson, Mark. *The Senses of Touch: Haptics, Affects and Technologies*. Berg, 2007.

Pedersen, Morten Axel, Kristoffer Albris, and Nick Seaver. "The Political Economy of Attention." *Annual Review of Anthropology* 50, no. 1 (2021): 309–25. https://doi.org/10.1146/annurev-anthro-101819-110356.

Potluri, Venkatesh, Tadashi E Grindeland, Jon E. Froehlich, and Jennifer Mankoff. "Examining Visual Semantic Understanding in Blind and Low-Vision Technology Users." In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–14.

35. New York, NY, USA: Association for Computing Machinery, 2021. https://doi.org/10.1145/3411764.3445040.

Pun, Karwai. "Dos and Don'ts on Designing for Accessibility - Accessibility in Government." Accessed May 20, 2021. https://accessibility.blog.gov.uk/2016/09/02/dos-and-donts-on-designing-for-accessibility/.

Rider, Gertrude T. *Braille Transcribing: A Manual*. Washington, D.C.: The American National Red Cross, 1925.

Schepers, Doug. "Why Accessibility Is at the Heart of Data Visualization." Medium, May 21, 2020. https://medium.com/nightingale/accessibility-is-at-the-heart-of-data-visualization-64a38d6c505b.

Sewell, Harry P. Method and apparatus for forming raised characters and lines. United States US2616198A, filed June 23, 1949, and issued November 4, 1952. https://patents.google.com/patent/US2616198A/en.

Sheppard, Linda, and Frances K. Aldrich. "Tactile Graphics in School Education: Perspectives from Teachers." *British Journal of Visual Impairment* 19, no. 3 (September 1, 2001): 93–97. https://doi.org/10.1177/026461960101900303.

Simply Accessible team. "7 Solutions for Creating More Accessible SVGs." *Simply Accessible* (blog), March 16, 2017. https://simplyaccessible.com/article/7-solutions-svgs/.

Sivek, Susan Currie. "Data Visualization and Accessibility: Three Recommended Reads and Top Tips." Medium, January 19, 2021. https://towardsdatascience.com/data-visualization-and-accessibility-three-recommended-reads-and-top-tips-9c5e862b464e.

Sorge, Volker, Charles Chen, T. V. Raman, and David Tseng. "Towards Making Mathematics a First Class Citizen in General Screen Readers." In *Proceedings of the 11th Web for All Conference*, 1–10. 40. New York, NY, USA: Association for Computing Machinery, 2014. https://doi.org/10.1145/2596695.2596700.

Soueidan, Sara. "Case Study: Implementing Accessible Data Charts for the Khan Academy 2018 Annual Report." Accessed May 20, 2021. https://sarasoueidan.com/blog/accessible-data-charts-for-khan-academy-2018-annual-report/.

Stoiber, Christina, Florian Grassinger, Margit Pohl, Holger Stitz, Marc Streit, and Wolfgang Aigner. "Visualization Onboarding: Learning How to Read and Use Visualizations." Preprint. Open Science Framework, August 15, 2019. https://doi.org/10.31219/osf.io/c38ab.

Stone, Brian W., Donovan Kay, and Anthony Reynolds. "Teaching Visually Impaired College Students in Introductory Statistics." *Journal of Statistics Education* 27, no. 3 (September 2, 2019): 225–37. https://doi.org/10.1080/10691898.2019.1677199.

Sullivan, J. E. "DOTSYS III: A Portable Braille Translator." In *Proceedings of the ACM Annual Conference*, 398–403. Atlanta, Georgia, USA: ACM, 1973.

Wang, Jing, Jinhui Tang, and Jiebo Luo. "Multimodal Attention with Image Text Spatial Relationship for OCR-Based Image Captioning." In *Proceedings of the 28th ACM International Conference on Multimedia*, 4337–45. Seattle WA USA: ACM, 2020. https://doi.org/10.1145/3394171.3413753.

Watson, Léonie. "Accessible SVG Flowcharts." Tink - Léonie Watson - On technology, food & life in the digital age. Accessed July 23, 2021. https://tink.uk/accessible-svg-flowcharts/.

———. "Accessible SVG Line Graphs." Tink - Léonie Watson - On technology, food & life in the digital age. Accessed July 23, 2021. https://tink.uk/accessible-svg-line-graphs/.

———. "Accessible SVG Tables." Tink - Léonie Watson - On technology, food & life in the digital age. Accessed July 23, 2021. https://tink.uk/accessible-svg-tables/.

Weninger, Markus, Gerald Ortner, Tobias Hahn, Olaf Druemmer, and Klaus Miesenberger. "ASVG Accessible Scalable Vector Graphics: Intention Trees To Make Charts More Accessible And Usable." *Journal of Assistive Technologies* 9 (November 10, 2015): 239–46. https://doi.org/10.1108/JAT-10-2015-0124.

Wijk, Jarke J. Van. "The Value of Visualization." In *IEEE Transactions on Visualization and Computer Graphics*, 12:79–86, 2006.

Williams, Sarah. *Data Action: Using Data for Public Good*. Cambridge, Massachusetts: The MIT Press, 2020.

Yu, Wai, Ramesh Ramloll, and Stephen Brewster. "Haptic Graphs for Blind Computer Users." Edited by Stephen Brewster and Roderick Murray-Smith. *Lecture Notes in Computer Science* 2058 (2001): 41–51. https://doi.org/10.1007/3-540-44589-7_5.

Zajicek, Mary, Chris Powell, and Chris Reeves. "A Web Navigation Tool for the Blind." In *Proceedings of the Third International ACM Conference on Assistive Technologies*, 204–6. Assets '98. New York, NY, USA: Association for Computing Machinery, 1998. https://doi.org/10.1145/274497.274534.

Zhang, Zeyuan. "Tapsonic: One Dimensional Finger Mounted Multimodal Line Chart Reader." In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, 1–4. ASSETS '20. New York, NY, USA: Association for Computing Machinery, 2020. https://doi.org/10.1145/3373625.3417075.

Zou, Hong, and Jutta Treviranus. "ChartMaster: A Tool for Interacting with Stock Market Charts Using a Screen Reader." In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility*, 107–16. ASSETS '15. New York, NY, USA: Association for Computing Machinery, 2015. https://doi.org/10.1145/2700648.2809862.

Conclusion

This dissertation investigates the process of curating, cleaning, visualizing, circulating, and manipulating data to understand the persuasive force of visual information in multimodal media. From the history of haptic interfaces to the data practices of social media communities across the US and China, this thesis uses historical and ethnographic methods to understand how users of quantitative information encode norms about gender, ability, and race in data visualizations and search interfaces. This critical scholarship complements projects with engineering colleagues at CSAIL to build more inclusive data representation systems. Drawing on work in feminist technoscience, disability studies, and the history and anthropology of computing, this dissertation weaves together different forms of HCI research to ask what work can or should be done by data representations across computational media.

The first two chapters of this thesis investigate the spread of COVID-related data and memes across social media. In *Let's Go, Baby Forklift!*, we draw from social media studies, linguistic anthropology, and China studies to investigate a phenomenon we call "fandom governance," which describes how state-run media borrows from a culture of celebrity fandom and memes in order to consolidate political power. *Viral Visualizations* similarly explores the political power of visual media by investigating how COVID-related data visualizations represent an ongoing epistemological gap between COVID skeptics and the scientific establishment. This study about the role of COVID data in public life is paired with the latter two chapters on alternative forms of data visualization: these chapters describe best practices for creating projects on accessible visualization, and instantiate them through prototypes for interactive data representations for screen readers. The prototypes in this project create new methods of interacting with a visualization that make it possible for blind and low vision users to seamlessly move between different layers and selections of data.

Exploring screen readers as a way to interactively present data challenges the ocularcentrism of conventional visualization practices, and it is an exciting opportunity for media scholars and critical makers to productively interrogate the relationship between the senses and interface design. The methodological variation in this dissertation is also an exciting artifact of how varied HCI and STS research can be. In surveys of the discipline by computer scientists, books like Rogers' (2012) *HCI Theory: Classical, Modern, and*

Contemporary and Erickson and McDonald's (2007) *HCI Remixed* tie the field's early developments with cognitive psychology and human factors research with additional interventions from ethnographers, computer scientists, and mechanical engineers. Given this breadth, I must return to the opening question of this dissertation — if HCI shape-shifts to include all manners of research, how can STS scholars understand HCI as a field and the kinds of knowledge production that it values? Historically, how has the preoccupation with efficiency, access, or scale in HCI research changed the imagined users and possible designs within a research project?

From perceptual psychology to design, and from anthropology to computer science, HCI in the last 10 years seems marked more by a common *interest* in the intersection between people and technology rather than by a common *method*. This kind of disciplinary cannibalism certainly has its perils; take one look at the overwhelming annual CHI conference and one might be hard-pressed to find a disciplinary through-line between ethnographic work on the virtual world *Second Life* and a systems building paper that describes the design of a new data visualization interface. Indeed, given the sundry nature of methods training across HCI programs — where certain programs tilt more to one end of the qualitative or quantitative spectrum — it can be difficult to see how an ethnographer might be able to rigorously evaluate the scholarly contributions of a systems building paper, and vice versa. While more conservative scholars might lament this kind of eclecticism as a kind of "undisciplining" — or, less charitably, as confused and incoherent — it is equally possible to imagine this disciplinary diversity as a foundry of unexplored scholarly pathways.

Many kinds of interdisciplinary work (either as a result of disciplinary imperialism or more peaceful means of academic collaboration) have meaningfully addressed these concerns by developing a shared vocabulary. For some, this means using shared concepts to encourage conversations on a personal level, but this kind of common language is even more effective on an institutional level (e.g., to advance a department's interest to a university administration, or to helm a professional society's entry into new disciplinary waters). It is not enough to simply acknowledge that HCI is diverse and porous compared to its more traditional counterparts like computer science. Science studies scholar Peter Galison, for example, has long argued that scientific work is *by nature* disunified and undisciplined, and it is this "*dis*unification [emphasis his] that underpins its strength and stability" (Galison 1999, 137).

In his analysis of the disciplinary distinctions between physicists (e.g., theoreticians, experimentalists, and instrumentalists, many of whom may not even call themselves "physicists" per se), Galison develops a concept of a "trading zone" to show how these different groups — as eclectic in concern and method they might be — join together at points of shared interest. "Like two cultures, distinct by living near enough to trade, they can share some activities while diverging on many others," Galison says. "What is crucial is that in the highly local context of the trading zone, *despite* the differences in classification, significance, and standards of demonstration, the two groups can collaborate. [...] They can even both understand that the continuation of exchange is a prerequisite to the survival of the larger culture of which they are part (ibid., 146). Susan Leigh Star and James R. Griesemer's (1989) concept of "boundary objects" in zoology has been similarly influential in describing the heterogeneous nature of scientific work and the need for collaboration that is crucial to science as an intellectual enterprise. In their formulation, a boundary object is something that is "both plastic enough to adapt to local needs and constraints...yet robust enough to maintain a common identity across sites" (ibid., 393). We might even construe these papers themselves as boundary objects that exist in a trading zone. Star and Griesemer's article is widely cited in information science and CSCW literature as well *as* STS; Galison is widely read in STS as well as history and media studies. Both papers are used as ways of talking about the heterogeneity of knowledge production in different disciplines. In fact, these approaches not only acknowledge the porousness of many disciplines, but they argue that this interchange is a *foundational component* to how disciplines tick. This is not just true of disciplines like HCI that we might see as an "umbrella" discipline that takes many subfields in traditional departments and makes them into its own; it is also visible within traditional disciplines like psychology or anthropology. There is little in common methodologically between perceptual and social psychologists, or between forensic and cultural anthropologists; however, by virtue of being in the same field, they share common questions — but different approaches to answering them.

Thus, it might be more instructive to establish what the common *questions* or *concepts* in HCI are, rather than quibble over whether there *should* be one. HCI has already been used as a useful term for researchers searching for a wide umbrella of academic interlocutors with whom to engage, but invoking that term also comes with its own historical baggage. At its outset, some of the first papers in flagship HCI journals deepen conceptual frameworks in computer science — itself a nascent field in the mid-20th century that might otherwise be called mathematics or electrical engineering — by articulating how these computational systems might be used in everyday life. In J.C.R. Licklider's "Man-Computer Symbiosis," for example, he articulates a future where there will be an "expected development in cooperative interaction between men and electronic computers" (Licklider 1960, 4) and similar luminaries like Alan Kay write of speculative futures where children and adults alike learn, work, and play primarily with digital computers (Kay 1972). That HCI's disciplinary origins come from discussions about the convergence of design and engineering à la Herbert Simon (1988) — alongside these speculative articles where human and computational work are becoming the site of shared intellectual inquiry — is often reflected in the kinds of work that scholarly HCI venues primarily value (i.e., systems building papers that describe future computer assisted work). Even more than that, the historical conversion of design into an engineering discipline — a process that was often hastened by industrial research organizations like SRI (formerly the Stanford Research Institute) — also foreshadows the larger scale of academic-industrial collaboration that is more prominent in HCI than perhaps any other social scientific or engineering discipline. The development of UX/UI research arms, along with the creation of roles like “product designer” as one that primarily requires an engineering degree with an artistic bent, is emblematic of the kinds of interdisciplinary (and extra-academic) collaborations that HCI has been able to foster. Understanding or problematizing these origin stories is important for talking about disciplinary formation because it is these stories that often circumscribe both what is possible and what is valued as valid forms of knowledge production.

These stories do not themselves clearly demarcate what HCI has been (or will be). However, framing what a discipline is around *common questions* therefore requires a normative sense of (1) what values are important in scholarly work, and (2) what kinds of questions can be important. Disciplinary training in a PhD is in large part gaining the tacit knowledge to understand the types of questions that are worth asking. For some, like Miriah Meyer or Paul Dourish, it is to create larger conceptual frameworks or methods to understand how design research can proceed (Sedlmair, Meyer, and Munzner 2012). For others, it can be to create computational visualization systems (e.g., Satyanarayan et al. 2017). To say simply that human-computer interaction encompasses anything to do with the intersection between people and computers seems too capacious, but this is a criticism that can be similarly levied at “traditional” disciplines like anthropology, which might proclaim to “study human culture.” Perhaps an alternative way to rigorously — but generously — scope what HCI as a discipline can do is to say that HCI researchers not only build technological artifacts, but they might also study *how* or *why* people use these artifacts in the first place. In this sense, scholars might be able to understand this field as a domain where some researchers focus both on systems of human or computational behavior, but they also examine the broader social and political contexts in which technological artifacts might be deployed (and by whom). This wide definition of what HCI can be is promising — it is the mode of doing scholarship that has encouraged the diversity of submissions at CHI — but it also attempts to more tightly scope the design space such that an individual article might be able to speak specifically to some components of that definition but not others.

This is not to throw all disciplinary boundaries to the wind and advocate for a chaotic world of anti-disciplinary thinking; most academics live in a world that enforces strict disciplinary boundaries at every turn: in qualifying examinations, in journal article submissions, and, perhaps most saliently, in tenure review. What is therefore valuable about this kind of interdisciplinary approach is the opportunity and the ability to reinterpret different kinds of knowledge from one discipline to another. Economist George Stigler (of “imperial science” fame) attributes economics’ success at conquering other disciplines to the abstractness and generalizability of economic theory (Stigler 1984); so too can these concepts and design frameworks be translated between qualitative and quantitative ways of knowing in HCI. Placing an emphasis in the design of computational systems not just on its constitutive parts but also on the people who use and design them is a first start — understanding how these different kinds of knowledge interact with each other is (hopefully) the future of this field.

Towards deeper intersections of sociotechnical work

A central tension between qualitative and technical areas of HCI is the epistemological role of **abstraction** and **generalizability**. In broad terms, one could argue that systems building in computer science deals primarily with abstraction and generality, whereas the humanities and social sciences prize close reading and champion specificity (e.g., this piece of knowledge *can only* be applied to *a specific group of people* within a *very specific time frame*; a historian would not believe that this particular phenomenon could be replicated across time, space, or culture). However, the impulse to generalize is not limited to the engineering sciences: theoretical research in the humanities and social sciences has also used granular data (e.g., ethnography) in order to make theoretical contributions that generalize about the nature of culture, the economy, or social organization (e.g., anthropological theories about embodiment). The impulse to generalize in technical HCI research follows similar patterns, where insights from wizard-of-oz user studies are used to inform broader guidelines about how to better design computational systems. However, these different types of generalized (and generalizable) knowledge are still validated in different ways: where quantitative research often aspires to be **reproducible,** the ultimate goal of many qualitative studies is to be **transferable**. There is no expectation, for example, that an ethnographer will be able to replicate the immersive field work of another scholar, but there is an expectation that one researcher's findings would make sense *by juxtaposition.* In both cases, however, understanding how knowledge is construed as legitimate or otherwise worth pursuing — and the ends to which it can be used — is useful for achieving a deeper intersection of sociotechnical work.

These kinds of binary analytical frames (see table 1) are pervasive throughout qualitative and quantitative research, and they are certainly not unique to HCI. Historians and anthropologists toggle between the local and the global, between micro- and macro-level events, and between observations and theory. Additionally, scientists and engineers often problematize the false dichotomy between pure and applied research to generate new developments in theory and methods (e.g., Stokes 1997). Many of these binaries seem to be largely a question of **scope** — it can be easy to simplify these distinctions as simply the analytical jump between "small" observations to "big" theories. These kinds of typologies are useful precisely because they can be organizing heuristics for the *data* we collect and for the *analytical frames* that we use to generate new knowledge.

Table 1

Qualitative	↔	*Quantitative*
Specificity	↔	Generalizability
Case study	↔	Abstraction
Descriptive	↔	Analytical
Deconstructive	↔	Constructive

In fact, many of the tensions between the qualitative and technical HCI researchers often stem from disagreements about how to make this analytical jump, or they can arise from misunderstanding if a particular intervention affects something on an individual, organizational, or social level. Take, for example, an anthropologist conducting an ethnography of a *particular* village in Papua New Guinea (case study), who then uses their observations to make a more general argument (an abstraction!) about the nature of gifts and human sociality. Many scientists and engineers can be suspicious of these kinds of methods, and for good reason: beyond the lack of replicability, and it certainly grates against the expectation that human behavior should be studied using a systematic sampling of a representative population. However, the same anthropologist might have misgivings about a computer scientist who randomly selects study participants throughout a larger pool to measure their reaction time with a certain piece of technology or to identify their political views. (I am not necessarily arguing whether or not these distinctions are warranted.) However, identifying them within different scales gives me an analytical framework to proceed with interdisciplinary collaboration that transcends distinctions between the "qualitative" and the "quantitative."

Indeed, many criticisms levied by social scientists of computational systems are that technical researchers have flattened the human experience into a few quantifiable dimensions and generalized without accounting for social context. Where a social scientist might criticize a computer scientist working on content moderation for because their tool does not meaningfully account for social factors like a power differential between users, the computer scientist might also respond that their technical interventions are not necessarily intended to, nor are they able to, change the ways that individual people behave. How, then, can we best resolve this tension between specificity and generality? Refracting research questions in HCI through the binary distinctions I present in figure 1 can be useful because we can map a project's data collection and analysis onto these frameworks in order to find new affordances (e.g., ethnographic data is specific but the theory created from it might be generalizable; similarly, we might try to take this kind of anthropological work — which is often deconstructive — and translate it into constructive guidelines for design). Bridging this divide may also require researchers to think about the different ways that the knowledge we produce can be **informative, predictive,** or **generative** (Beaudouin-Lafon 2000). In Beaudouin-Lafon's model, the most powerful research is *generative,* such that it helps researchers envision new possibilities within a computational system; formalisms are helpful for understanding, guiding, and building new systems. Connecting these kinds of technical insights to qualitative research might then require what Yvonne Rogers (2005) calls a hybrid approach: the ability to be *generative* (i.e., provide design dimensions to inform the design and selection of interactive representations) and *formative* (e.g., a set of easy-to-use concepts for discussion a design) requires *rich descriptive accounts* that explain user behavior, and they also depend on *analytic frameworks* (e.g., high-level conceptual tools for identifying problems). Most importantly, the binary frameworks as described in table 1 give researchers an outline for understanding the limits and possibilities of *deconstructive analysis* (i.e., most, if not all, of STS literature on AI ethics) with *constructive* *analysis* (e.g., systems building that emphasizes fairness). This balance is what allows HCI researchers to productively identify the possibilities of collaboration.

Allow me to illustrate the affordances and constraints of thinking this way by showing more concretely how design studies and ethnography can inform systems design (and vice versa). Design study methodology (Sedlmair, Meyer, and Munzner 2012) yields three kinds of research contributions: problem characterization, validated visualization design, and reflection. Ethnography yields insight into a culture via long-term immersion, which can show HCI researchers how the organization of social settings can affect the kinds of systems which can (or should) be deployed there. Ethnography and design studies therefore have a critical role to play in interactive systems design because they can shape overarching research questions and uncover the constraints or opportunities in the way that a particular project has been framed.

While this can be the practical basis by which ethnography contributes to the design of a new system, Paul Dourish (2014) argues that ethnography *also* creates a kind of abstraction on which scientific generalization relies. Methodologically, it generalizes *through juxtaposition* — using "contradistinction, comparison, sequentiality...and other ways of patterning across multiple observations. This form of ethnographic juxtaposition does not itself truck in abstractions, but it extends itself beyond the circumstances of specific observation" (ibid.,13). These qualitative studies can therefore be used to shape the infrastructure and features of a particular technical system; instead of thinking about how systems can be disseminated to users *after the fact* (e.g., Ames 2019), HCI researchers can better understand the social and cultural ecosystems in which technological projects are intended to intervene.

The same logic — but in the opposite direction — may be even easier to illustrate. Examples of cultural and social critique abound of computational systems where the author themselves have very little technical expertise; in these examples, it is possible to completely misinterpret emic categories and change the meaning of what an interlocutor might have said — or worse, what the system can or cannot do. The deployment of AI systems are terrifying for a whole host of reasons, but usually it is because these kinds of systems are highly elementary sorting algorithms but made out to be surgically precise machines, which makes its presence in criminal justice proceedings all the more dystopic. Qualitative studies of how computational systems are used and suffused in everyday activities are immeasurably improved by having an incisive understanding of how or why certain types of systems are built (which is but one of the reasons why I am currently a member of this group). This kind of expertise, when refracted through a prism of social theory, disability studies, and science studies, becomes much more vivid data to analyze precisely because it ties together so many different dimensions of how people interact with computational systems.

In sum, I believe that the convergence of social and technical work can be facilitated through a critical re-examination of the epistemological role of abstraction and generalizability. Both kinds of work engage in abstraction — whether through the production of anthropological theory, or through formalisms like direct manipulation. So too do they traffic in different types of generalizability: ethnography generalizes through juxtaposition, and technical work generalizes often to create new design frameworks. However, by working at both of these registers, I believe that HCI can mobilize social insights in order to inform current design, and that technical expertise can be used to further excavate the ways that race, gender, and ability have been encoded in the design of human-machine interfaces.

Intellectual contributions

By thinking about the social life of data, this dissertation's first intellectual contribution will be to the rich literature in critical disability studies. Scholars in this field have written extensively on the relationship between the body, technology, and disability and the ways that the design of technological artifacts reifies existing notions about what is considered "normal," often to horrifying results (M. Mills 2011; Hamraie 2017; Hendren 2017; Davis 2013; Thomson 1996). Mara Mills, for example, has shown how engineers used deafness as a pretext for developments in telephony, where inventors often abandoned their deaf students after initial trials were completed, or they used deafness as a way of publicizing their flashy new technologies (2010). This story has important resonances today as technologists extol the potential for tactile interfaces to enhance the accessibility of data visualizations; the latter two chapters in this dissertation leverage this work to think about how to conscientiously design these prototypes and show how they create additional opportunities for thinking about visualization structures.

Conceptually, this project also draws together critical histories of data science (D'Ignazio and Klein 2016; Aronova, Oertzen, and Sepkoski 2017; M. L. Jones 2018). In sociologist David Ribes' words, STS scholars should critically engage with data science, a field which "[seeks] to touch virtually every science and newly bind spheres of academy, industry, and the state" (2019). By asking how these data become physically instantiated through touch, I refract this literature on visualization (Daston and Galison 2007; Dumit and Burri 2008; Coopmans 2014) in new ways that emphasize how a scientist's sensorium can mediate her understanding and interpretation of scientific data, where the senses of sound and touch are active participants in the process of scientific discovery (Helmreich 2007; Vertesi 2012; Myers 2015). More broadly, my research also follows a long line of STS scholars who show how social scientists can work collaboratively and productively with scientists and engineers (Sengers et al. 2005; Eglash et al. 2006; Wylie et al. 2014). In so doing, my project contributes to the vibrant literature from humanists who interrogate the assumptions, biases, and histories of modern computing practices that deal specifically with non-Western science (Medina 2011; Mullaney 2017) and disability (Petrick 2015).

Finally, I also contribute to literature by computer scientists who use qualitative research methods and insights in their studies of interactive systems design (Dourish 2004; Taylor 2009; Rosner and Fox 2016; Howard and Irani 2019). By writing for both communities, my research mobilizes STS theory into practice. I am critical of the social values and practices that become embedded in technology design, but I draw on that history to then develop collaborative methods for making future technologies that are anti-racist, feminist, and inclusive (Guston 2014). As Bronet and Layne write: “We do not want to sit back and offer post facto critiques of new technologies, but want to intervene proactively to influence design. One obvious way to do this is to equip the next generations of designers to work towards this goal” (2010, 179). In this sense, this work contributes to growing conversations about how qualitative studies of science and engineering can be used for critical making (Ratto 2011; Hamraie 2018) and reflective design (Sengers et al. 2005).

References for Conclusion

Ames, Morgan G. 2019. *The Charisma Machine: The Life, Death, and Legacy of One Laptop per Child*. Infrastructures. Cambridge, MA, USA: MIT Press.

Aronova, Elena, Christine von Oertzen, and David Sepkoski. 2017. “Introduction: Historicizing Big Data.” *Osiris* 32 (1): 1–17. https://doi.org/10.1086/693399.

Beaudouin-Lafon, Michel. 2000. “Instrumental Interaction: An Interaction Model for Designing Post-WIMP User Interfaces.” In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 446–53. CHI ’00. New York, NY, USA: ACM. https://doi.org/10.1145/332040.332473.

Bronet, Frances, and Linda L. Layne. 2010. “Teaching Feminist Technology Design.” In *Feminist Technology*, edited by Sharra Louise Vostral, Kate Boyer, and Linda L. Layne. University of Illinois Press.

Coopmans, Catelijne. 2014. “Visual Analytics as Artful Revelation.” In *Representation in Scientific Practice Revisited*, edited by Catelijne Coopman, Janet Vertesi, Michaeland Lynch, and Steve Woolgar, 37–58. The MIT Press. https://doi.org/10.7551/mitpress/9780262525381.003.0003.

Daston, Lorraine, and Peter Galison. 2007. *Objectivity*. Cambridge, MA: Zone Books.

Davis, Lennard J. 2013. *The End of Normal: Identity in a Biocultural Era*. Ann Arbor: University of Michigan Press.

D’Ignazio, Catherine, and Lauren F. Klein. 2020. *Data Feminism*. Strong Ideas Series. Cambridge, Massachusetts: The MIT Press.

Dourish, Paul. 2004. “What We Talk about When We Talk about Context.” *Personal and Ubiquitous Computing* 8 (1): 19–30. https://doi.org/10.1007/s00779-003-0253-8.

———. 2014. “Reading and Interpreting Ethnography.” In *Ways of Knowing in HCI*, edited by Judith S. Olson and Wendy Kellogg, 1–24. New York: Springer.

Dumit, Joseph, and Regula Valérie Burri. 2008. “Social Studies of Scientific Imaging and Visualization.” In *The Handbook of Science and Technology Studies*, edited by Edward J. Hackett, Olga Amsterdamska, Judy Wajcman, and Michael Lynch, 297–317. Cambridge, MA: MIT Press.

Eglash, Ron, Audrey Bennett, Casey O’Donnell, Sybillyn Jennings, and Margaret Cintorino. 2006. “Culturally Situated Design Tools: Ethnocomputing from Field Site to Classroom.” *American Anthropologist* 108 (2): 347–62. https://doi.org/10.1525/aa.2006.108.2.347.

Erickson, Thomas, and David W. McDonald, eds. 2007. *HCI Remixed: Reflections on Works That Have Influenced the HCI Community*. Cambridge, MA, USA: MIT Press.

Galison, Peter. 1999. “Trading Zone: Coordinating Action and Belief (1998 Abridgment).” In *The Science Studies Reader*, edited by Mario Biagioli, 137–60. London; New York: Routledge.

Guston, David H. 2014. “Understanding ‘Anticipatory Governance.’” *Social Studies of Science* 44 (2): 218–42. https://doi.org/10.1177/0306312713508669.

Hamraie, Aimi. 2017. *Building Access: Universal Design and the Politics of Disability*. Minneapolis, MN: University of Minnesota Press.

———. 2018. “Mapping Access: Digital Humanities, Disability Justice, and Sociospatial Practice.” *American Quarterly* 70 (3): 455–82. https://doi.org/10.1353/aq.2018.0031.

Helmreich, Stefan. 2007. “An Anthropologist Underwater: Immersive Soundscapes, Submarine Cyborgs, and Transductive Ethnography.” *American Ethnologist* 34 (4): 621–41. https://doi.org/10.1525/ae.2007.34.4.621.

Hendren, Sara. 2017. “All Technology Is Assistive: Six Design Rules on Disability.” In *Making Things and Drawing Boundaries*, edited by Jentery Sayers, 139–46. Experiments in the Digital Humanities. University of Minnesota Press. https://doi.org/10.5749/j.ctt1pwt6wq.18.

Howard, Dorothy, and Lilly Irani. 2019. “Ways of Knowing When Research Subjects Care.” In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI ’19*, 1–16. Glasgow, Scotland: ACM Press. https://doi.org/10.1145/3290605.3300327.

Jones, Matthew L. 2018. “How We Became Instrumentalists (Again): Data Positivism since World War II.” *Historical Studies in the Natural Sciences* 48 (5): 673–84. https://doi.org/10.1525/hsns.2018.48.5.673.

Kay, Alan C. 1972. “A Personal Computer for Children of All Ages.” In *Proceedings of the ACM Annual Conference - Volume 1*. ACM ’72. New York, NY, USA: ACM. https://doi.org/10.1145/800193.1971922.

Licklider, J.C.R. 1960. “Man-Computer Symbiosis.” *IEEE Transactions on Human Factors in Electronics*, March, 4–11.

Medina, Eden. 2011. *Cybernetic Revolutionaries: Technology and Politics in Allende’s Chile*. Cambridge, MA, USA: MIT Press.

Mills, M. 2011. “On Disability and Cybernetics: Helen Keller, Norbert Wiener, and the Hearing Glove.” *Differences* 22 (2–3): 74–111. https://doi.org/10.1215/10407391-1428852.

Mills, Mara. 2010. "Deaf Jam: From Inscription to Reproduction to Information." *Social Text* 28 (102): 35–58. https://doi.org/10.1215/01642472-2009-059.

Mullaney, Thomas S. 2017. *The Chinese Typewriter: A History*. 1 edition. Cambridge, MA: MIT Press.

Myers, Natasha. 2015. *Rendering Life Molecular: Models, Modelers, and Excitable Matter*. Experimental Futures : Technological Lives, Scientific Arts, Anthropological Voices. Durham: Duke University Press.

Olson, Judith S., and Wendy Kellogg, eds. 2014. *Ways of Knowing in HCI*. New York: Springer.

Petrick, Elizabeth R. 2015. *Making Computers Accessible: Disability Rights and Digital Technology*. Baltimore, MD: Johns Hopkins University Press.

Ratto, Matt. 2011. "Critical Making: Conceptual and Material Studies in Technology and Social Life." *The Information Society* 27 (4): 252–60. https://doi.org/10.1080/01972243.2011.583819.

Ribes, David. 2019. "STS, Meet Data Science, Once Again." *Science, Technology, & Human Values* 44 (3): 514–39. https://doi.org/10.1177/0162243918798899.

Rogers, Yvonne. 2005. "New Theoretical Approaches for Human-Computer Interaction." *Annual Review of Information Science and Technology* 38 (1): 87–143. https://doi.org/10.1002/aris.1440380103.

———. 2012. *HCI Theory: Classical, Modern, and Contemporary*. Synthesis Lectures on Human-Centered Informatics 14. San Rafael, CA: Morgan & Claypool.

Rosner, Daniela K, and Sarah E Fox. 2016. "Legacies of Craft and the Centrality of Failure in a Mother-Operated Hackerspace." *New Media & Society* 18 (4): 558–80. https://doi.org/10.1177/1461444816629468.

Satyanarayan, Arvind, Dominik Moritz, Kanit Wongsuphasawat, and Jeffrey Heer. 2017. "Vega-Lite: A Grammar of Interactive Graphics." In *IEEE Transactions on Visualization & Computer Graphics (Proc. IEEE VIS)*. https://doi.org/10.1109/TVCG.2016.2599030.

Sedlmair, Michael, Miriah Meyer, and Tamara Munzner. 2012. "Design Study Methodology: Reflections from the Trenches and the Stacks." *IEEE Transactions on Visualization and Computer Graphics* 18 (12): 2431–40. https://doi.org/10.1109/TVCG.2012.213.

Sengers, Phoebe, Kirsten Boehner, Shay David, and Joseph "Jofish" Kaye. 2005. "Reflective Design." In *Proceedings of the 4th Decennial Conference on Critical Computing: Between Sense and Sensibility*, 49–58. CC '05. New York, NY, USA: ACM. https://doi.org/10.1145/1094562.1094569.

Simon, Herbert A. 1988. "The Science of Design: Creating the Artificial." *Design Issues* 4 (1/2): 67. https://doi.org/10.2307/1511391.

Star, Susan Leigh, and James R. Griesemer. 1989. "Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39." *Social Studies of Science* 19 (3): 387–420. https://doi.org/10.1177/030631289019003001.

Stigler, George J. 1984. "Economics: The Imperial Science?" *The Scandinavian Journal of Economics* 86 (3): 301–13. https://doi.org/10.2307/3439864.

Taylor, Alex S. 2009. "Ethnography in Ubiquitous Computing." In *Ubiquitous Computing Fundamentals*, edited by John Krumm, 203–37. Chapman and Hall/CRC.

Thomson, Rosemarie Garland. 1996. *Extraordinary Bodies: Figuring Physical Disability in American Culture and Literature*. 1st edition. New York: Columbia University Press.

Vertesi, Janet. 2012. "Seeing like a Rover: Visualization, Embodiment, and Interaction on the Mars Exploration Rover Mission." *Social Studies of Science* 42 (3): 393–414. https://doi.org/10.1177/0306312712444645.

Wylie, Sara Ann, Kirk Jalbert, Shannon Dosemagen, and Matt Ratto. 2014. "Institutions for Civic Technoscience: How Critical Making Is Transforming Environmental Research." *The Information Society* 30 (2): 116–26. https://doi.org/10.1080/01972243.2014.875783.