Physical Rendering Processes for More Graspable Extended Reality Data Visualizations

A dissertation submitted to the faculty of the University of Minnesota by

Bridger Herman
Advised by Dr. Daniel F. Keefe

In partial fulfillment for the degree of Doctor of Philosophy in Computer Science.

December 2023

Acknowledgements

I want to acknowledge that whatever words I write here will never be enough to recognize all of the individuals who have supported, encouraged, and made it possible for me to be where I am today. To everyone who has had a role in helping me become the person I am today and to those whose guidance and kindness were instrumental in helping me finish this dissertation, including all those mentioned below: thank you. To begin, I want to express my most sincere gratitude for my family. To my parents, who have never stopped believing that I can be my very best self. Your support through my life, both emotional and practical, has been an integral part of my academic journey. I may be a "B student", but that's okay because my work doesn't need to be perfect to be good, or to make a difference. To my partner, Emily, whose kindness and support every day of my academic career has made it possible for me to get this far (especially over the past year): your work does not go unnoticed. Truly, I value your assistance as my auxiliary brain when I inevitably forget to do basic self-care, I sincerely appreciate our teamwork relationship, and I look forward to being a full half of the team once again.

To my advisor, Dan Keefe: thank you for your mentorship in all aspects of academia during my entire time in the lab. Between teaching me how to be a more specific writer, developing course material in various classes, all the system-building and programming time, and helping me navigate the challenges of being graduate student (especially the hard pivots in research directions necessitated by the pandemic), your guidance has been essential. And most importantly, thank you for always keeping in mind that your students are human, too. My deepest gratitude to you for being a mentor and a role model in so many ways over the past seven years of my life, and through the five years of the "joyful struggle" of my Ph.D.

I also want to thank my committee members, Victoria Interrante, Cullen D. Jackson, and Lana Yarosh, whose expertise and discussions have made the work in this dissertation – especially the study in Chapter 6 – so much more than it was to begin with, and in the process of doing so, have helped me discover more about myself as a researcher. Beyond my committee, I would like to acknowledge Stephen J. Guy and Evan Suma Rosenberg, who never let me forget that computer graphics can (and should) be fun, Greg Abram, for always being down to build technical prototypes, and Francesca Samsel, who always reminds us *why* we're building these prototypes and always finds new, creative uses for the things we build. And to all my collaborators and co-authors over the years: thank you for the opportunity to work together and for opening my mind to different disciplines of research and beyond.

I'd like to thank current and past members of the Interactive Visualization lab that I've worked with. To my current fellow Ph.D. lab mates, Sean Dorr and Kiet Tran, thanks for always being there to discuss anything, from academics, to teaching, to design, to life beyond academia, I appreciate every conversation I am privileged to have with you both. And to Alex Haley, though our paths through our respective Ph.D.s have been quite different, I always appreciate the perspectives you bring to whatever conversation we happen to be having. To the Ph.D. students who mentored me early in my early days of research, Seth Johnson, Dan Orban, Volcano Kim, and Jung Who Nam: your support, mentorship, and friendship through the early days of my research was truly instrumental. To Morgan Turner: thanks for always reminding me of the "why" of research. And to the current and past undergraduate and Masters students I've had the honor of working with: keep doing your best work, and thank you for the opportunity to work with you.

Lastly, I'd like to express my gratitude for my friends. To my current and past Minnesota Brass family: thank you for always reminding me that there's more to life than academics. To the graphics squad, Nikki Kyllonen and Liam Tyler: thank you both for your encouragement, friendship, and shenanigans in early grad school. To Arie Van Antwerp, thanks for giving me an excuse to come visit, which truly set the direction for my future academic career. And to all the others who have taken the time over the past five years to keep reaching out to stay in touch, despite my seemingly continual response of "sorry, I can't – I'm writing."

For all the educators in my life.
Your belief in me is what made this work possible.

Abstract

This dissertation examines new approaches to extended reality visualization that use physical rendering processes to make data more graspable: both more comprehensible and more embodied. Current extended reality visualizations enable new data insights and better comprehension of spatial data, but they lack the physical embodiment and tactility that has been crucial to past scientific breakthroughs. Current data physicalizations have this tactility yet lack in the interactivity required for today's scientific sensemaking processes. In this dissertation, I explore new ways to combine the embodiment of data physicalizations and the comprehensibility of extended reality visualizations to make data more graspable using physical rendering processes, which make it possible to create physical objects from digital data and enable physical elements from the real world to be used in digital visualizations.

The dissertation presents four main contributions supported by physical rendering processes that span the extended reality continuum. To begin, I present a systematic design exploration of the first 3D printed spatial data physicalizations encoding scalar data as glyphs on a 3D surface using traditional, forward physical rendering processes. Then, I introduce a new, *inverse* physical rendering process, a software architecture and user interface enabling the design of multivariate 3D spatial data visualizations with handcrafted physical media. Combining the forward and inverse physical rendering processes from the first two contributions, I penultimately present a new approach to querying 3D spatial data using multi-touch input directly on a data physicalization. Finally, I present the design and results of the first empirical study on the effectiveness of data physicalizations for spatial data analysis tasks when compared with state-of-the-art virtual reality and 2D geographic information science visualizations.

Multiple conclusions are drawn from the explorations into spatial data physicalizations and extended reality visualization. First, data physicalizations have the potential to be more comprehensible than digital visualizations when completing spatial data analysis tasks, and such benefits are most likely tied to the act of *physically touching* the data. Second, inverse physical rendering processes that encode digital data with elements from the physical world bring a new level of embodiment to digital extended reality visualizations, and they show many of the benefits of physicalization, potentially through *imagined touch*. Lastly, new approaches that tightly integrate physicalizations with digital interactivity lead to new opportunities for collaboration and engagement in scientific sensemaking processes.

Contents

List of Figures **vii**

List of Tables **viii**

1 Introduction **1**
- 1.1 Physical Rendering Processes 5
- 1.2 Towards More Graspable Data 6
- 1.3 Conceptual Framework and Thesis Statement 9
- 1.4 Key Contributions of this Dissertation 11
- 1.5 Overview of the Dissertation 11

2 Related Work **13**
- 2.1 Making Data More Comprehensible 14
- 2.2 Tangible and Embodied Interaction 18
- 2.3 Physical Rendering Processes 19

3 Designing Glyph-Based Spatial Data Physicalizations **23**
- 3.1 Introduction 24
- 3.2 Methodology 24
- 3.3 Designing Glyph-Based Spatial Data Physicalizations 27
- 3.4 Fabricating Glyph-Based Spatial Data Physicalizations 30
- 3.5 Results and Discussion 32
- 3.6 Conclusion 33

4 Creating Spatial Data Visualizations with Handcrafted Physical Media **34**
- 4.1 Introduction 35
- 4.2 Related Work 36
- 4.3 Design Requirements 40
- 4.4 Implementation 42
- 4.5 Results 52
- 4.6 Discussion 56
- 4.7 Conclusion 58

5 Multi-Touch Querying on Data Physicalizations in Immersive AR **60**
- 5.1 Introduction 61
- 5.2 Related Work 64
- 5.3 Hardware and Strategies for Addressing Technical Challenges 69
- 5.4 System and Interaction Techniques for Multi-Touch Querying on Data Physicalizations 72
- 5.5 Evaluation and Application 75
- 5.6 Discussion 80
- 5.7 Conclusion 82

6 Touching the Ground: Evaluating the Effectiveness of Data Physicalizations for Spatial Data Analysis Tasks **83**
6.1 Introduction 84
6.2 Related Work 86
6.3 Study Design 90
6.4 Results 103
6.5 Discussion 112
6.6 Limitations and Future Work 119
6.7 Conclusion 121

7 Conclusion **122**
7.1 Review of Contributions Towards Thesis Statement 123
7.2 Future Research 125
7.3 General Conclusions 126

Bibliography **128**

List of Figures

1.1 The Physical Rendering Process 5
1.2 Graspable data physicalization combined with interactive extended reality scientific visualization 8
1.3 Physical Rendering Processes for More Graspable XR Visualizations . . 9

3.1 Current design for glyph-based spatial data physicalizations 26
3.2 Design space for glyph-based physicalizations 27
3.3 Glyph shapes used in the design process 28
3.4 Fabrication of glyph-based physicalizations 31

4.1 The visualization architecture and user interface enables artists to encode multivariate spatial data with physical, real-world elements. 35
4.2 Comparison of computer graphics and printmaking processes 41
4.3 Construct a glyph data impression by applying data, a colormap, and a 3D scanned clay artifact 42
4.4 A glyphs data impression in the user interface 43
4.5 Visualization software architecture diagram 45
4.6 The design user interface 47
4.7 Comparison of applying glyphs vs. ribbons to a line 49
4.8 The custom colormap editor 50
4.9 Example visualization produced by the JSON schema 51
4.10 Examples created by artists with the design interface and architecture . 53

5.1 Multi-touch augmented data physicalization supporting volumetric data queries 61
5.2 Design and construction of the tracking base and augmented physicalization 68
5.3 Experiment setup to measure and minimize latency 70
5.4 Equipment setup with antarctic context and AR view of the physicalization 72
5.5 Seed streamlines by touching with one finger, create cutting planes with two fingers 73
5.6 Experimental setup for touch accuracy validation 75
5.7 Application of multi-touch augmented physicalization to climate science 78

6.1 Spatial data physicalization compared with 2D and virtual reality visualization 84
6.2 Visualization design for the study 92
6.3 Task overlay and tasks for the study 97
6.4 Study procedure 99
6.5 Experimental setup for study 102
6.6 Study results 104
6.7 Differences between Physical and other modalities 105

List of Tables

1.1 Organization of the dissertation 12

6.1 Elevation datasets for the user study 96
6.2 Touch and manipulation in the Physical modality 109
6.3 Positive factors 110
6.4 Negative factors 110
6.5 Strategies 111

chapter one

Introduction

1.1 Physical Rendering Processes 5
1.2 Towards More Graspable Data 6
1.3 Conceptual Framework and Thesis Statement 9
1.4 Key Contributions of this Dissertation 11
1.5 Overview of the Dissertation 11

For millennia, humans have used physical objects to help us make sense of the universe; from clay tokens representing trade and profit figures in Mesopotamia around 4000 BCE [Schmandt-Besserat, 2013] to Athenian citizens casting votes for political candidates using colored stones [Boegehold, 1963], these physical objects have facilitated a better understanding of the world by providing a graspable, tangible representation of numerical quantities, abstract data, and models of physical phenomena. The use of physical objects has accompanied advances in the scientific method and scientific discoveries developing new insights about complex concepts and phenomena, and communicating these discoveries with others. From the 17th century onwards, physical objects have been used to understand and teach the intricacies of human anatomy to provide continually better medical care [Riva et al., 2010]. In the late 1800s, physical objects were essential to the discovery of the thermodynamic model of state; the thermodynamcist Gibbs used these physical data representations to comprehend the nonlinear relationships between volume, energy, and entropy [Kriz, 2007]. In the last century, physical objects have been used in conjunction with advances in microscopy and imaging to better understand macromolecules like myoglobin, which stores oxygen within animal cells [Kendrew et al., 1958]. Physical data representations have remained a vital part of our scientific sensemaking processes throughout their history, due in part to the range of information that they convey through their geometric and material properties, and their innate capacity to support tactile exploration and feedback: as naturally as we interact with everyday objects in the real world, we can touch, hold, mark points with fingers, and rapidly move and rotate physical data objects to develop a better comprehension of the universe around us.

Today, *Scientific Visualization* based on computer graphics holds a similar value to modern data-intensive[1] scientific sensemaking processes that physical objects held prior to the computational revolution. Scientific visualization is a research field devoted to utilizing "*the broad bandwidth of the human sensory system in steering and interpreting the complex processes and simulations involving voluminous data sets across diverse scientific disciplines.*" [Nielson et al., 1997]. Generally, research in this field is advanced through close collaborations between domain scientists and computer scientists. Domain scientists have driving grand challenges in their disciplines (e.g., developing more sustainable energy solutions with biofuels, or understanding how Antarctic ice is melting), and computer scientists simultaneously advance the visualization field by developing novel interfaces, techniques, and approaches to making the domain scientists' data more comprehensible and address their grand challenges.

Current widespread scientific visualization software (e.g., ParaView [Ahrens et al.,

[1] Data-intensive computing, particularly scientific computing, frequently requires terabyte or petabyte sized datasets that require special hardware and software to work with.

2005], VisIt [Childs et al., 2012]) enables scientists to make queries into their data to answer questions about, and develop better comprehension of the phenomena they are studying: from seeding streamlines based on supercomputer simulations of ocean current flow data for understanding the melting pattern of polar ice caps to distributing glyphs based on nitrate levels in the ocean for evaluating the productivity of future macroalgae farms supporting more sustainable energy solutions, this software has been critical to understanding multivariate, time-varying 3D spatial data that are too large, dense and complicated to make sense of without computers[2]. Most importantly, scientific visualization software enables scientists to make queries in near real-time, providing rapid feedback which is used to explore, to test theories, and to better comprehend their data. Additionally, many scientific visualization software packages enable scientists to visualize multiple variables simultaneously to better understand the relationships and trends the data exhibit. Through the power of computer graphics, scientific visualizations make it possible to get information, and ultimately understanding, from the data and relationships presented by modern scientific challenges.

Data physicalization takes a decidedly different approach to addressing the needs of visualization stakeholders by encoding data with the geometry and material properties of physical objects [Jansen et al., 2015]. Though the term "data physicalization" is relatively new, it builds on ideas stemming from traditional physical sensemaking processes: physical objects innately support certain interactivity and active perception like touching, marking points, and manipulating, they enable us to depict and communicate complex spatial data in a way that directly maps to physical space, and they take advantage of the multisensory nature of our human perceptual systems by leveraging both visual *and* haptic senses and make it possible to use our embodied, human reference frame to make size and distance judgments. Physicalizations also enable true, real-world depth perception (as opposed to simulated depth in immersive digital displays), they make data analysis more accessible to low-vision and blind users, and they can better facilitate understanding and learning [Jansen et al., 2015].

The process of creating physical data representations is known as the *physical rendering process* in the data physicalization literature [Jansen and Dragicevic, 2013, Djavaherpour et al., 2021]. Current data physicalizations are most often created with physical rendering processes that employ digital fabrication technologies like 3D printing, laser cutting, and CNC milling, which are becoming increasingly cost-effective and easier to use with the maker revolution [Willis, 2018]. Yet, despite their decreasing cost, their possible benefits, and the historic precedent for their use, data physicalizations have not

[2] Multivariate = multiple quantities at each data point,
Time-varying = data contains multiple frames in time,
3D spatial data = data that inherently have a structure defined in 3D space.

been widely adopted in current scientific sensemaking processes.

Extended reality (XR) visualization is an active research area in computer science aimed at making current computer graphics visualizations more graspable with additional hardware and software to solve some key problems of "traditional" desktop computing based computer graphics visualizations, and is also known as *immersive* visualization. XR spans a continuum of display technologies ranging from entirely digital virtual reality (VR) displays, to augmented (AR) or mixed (MR) reality displays that combine digital and real-world physical elements, up to and including displays rooted in the physical world [Milgram and Kishino, 1994]. In scientific visualization, a key advantage of using extended reality is that it enables scientists to view their "3D data in 3D", use spatial input devices to interact with their data in ways that emulate everyday physical interactions, and even engage the senses of touch or hearing, via haptics or sound [Bryson, 1996].

Today, XR scientific visualizations have in many ways surpassed the capabilities of their previous counterparts, physical data representations. These visualizations enable all the benefits of the best digital visualization software necessitated by today's data-intensive scientific sensemaking processes; they make interactive querying into large datasets possible, and they can display 3D spatial data in 3D. In the last decade, interest in XR visualization has become widespread across the sciences, and examples of XR visualization usage can be found in geographic information science [Çöltekin et al., 2020], biomedical science [Venkatesan et al., 2021], and even archaeology [Bonis et al., 2022]. However, despite technological advancements and growing adoption in real-world use cases, digital XR visualizations remain inferior to earlier physical data representations in some important aspects: Digital XR visualizations are highly interactive, yet they leave out an important component of human perception and sensemaking: touch. Digital XR visualizations enable scientists to rapidly work with terabytes of data at a time, yet they often add layers of indirect input/output devices that distance users from their data. Digital XR visualizations are frequently used to represent real-world phenomena, yet they rarely incorporate real-world elements that would provide more channels in which to encode the increasingly complex multivariate spatial data required by today's science. Notwithstanding their current limitations, digital XR visualizations continue to be a growing part of scientific sensemaking processes, and practitioners in the field continue to research, create, and implement new, more effective methods to visualize scientific data.

1.1 Physical Rendering Processes

Contrary to the digital rendering processes used by computer graphics visualizations which are perceived by users with a 2D screen or XR display, *physical rendering processes* make data "*perceivable by bringing it into existence in the physical world*" [Jansen and Dragicevic, 2013]; these are the processes with which data physicalizations are created. There are multiple emerging perspectives to considering the physical rendering process, including the encodings and properties used [Hornecker et al., 2023], the design [Bae et al., 2022], and the fabrication [Djavaherpour et al., 2021] of data physicalizations. However, current perspectives of physical rendering processes are becoming insufficient to describe the increasingly intertwined nature of digital and physical visualizations. As they are currently defined, physical rendering processes explicitly take digital data and create physical, real-world representations of them, but *they do not consider using elements from the physical world in digital visualizations.* If data visualization practitioners are to truly leverage the possibilities enabled by encoding data with physical objects, the definition of physical rendering processes must be broadened.

I propose a new, expanded definition of the *physical rendering process* in the context of data visualization where data coexist between the physical world and digital world:

> *Physical rendering processes enable digital data to be visualized in the physical world, and they make it possible for elements from the physical world to be used in digital data visualizations.*

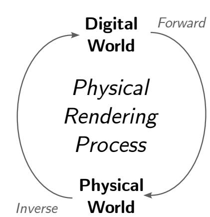

Figure 1.1

This new definition, shown in Figure 1.1, encompasses a future of data physicalization for scientific sensemaking processes where the digital and physical elements of a visualization are tightly integrated like the extended reality visualization in Figure 1.2 and those contained in Chapter 5. In these novel examples, the digital data are transformed into the physical world via data physicalizations, and elements from the physical world are used to encode data in digital extended reality visualizations.

The new definition is a superset of the original; it equivalently encompasses prior work including static physicalizations that are only rendered a single time through digital fabrication, active physicalizations which update and re-render their physical geometry based on interactions or changes to the data, and even augmented physicalizations that utilize both physical and digital components. The largest change from prior definitions comes into play when elements from the physical world are explicitly used to visualize digital data, as in the Artifact-Based Rendering technique [Johnson et al., 2019b], which make it possible for designers and artists to handcraft data encodings with physical media to use in extended reality visualizations. Ultimately, this definition anticipates the future of extended reality visualizations where the boundaries between the digital and physical are blurred: where scientists can directly touch a data physicalization to make queries into their 3D multivariate spatial datasets, and query results are displayed as digital extended reality visualizations blended with the physicalization; where artists collaborate with scientists to make the most comprehensible and engaging visualizations using materials directly from their studios; and where the benefits of *both* highly interactive digital visualizations and embodied physical visualizations are shared by all.

1.2 Towards More Graspable Data

Prior work in both extended reality visualization and data physicalization (by using physical rendering processes) has pursued the goal of making data more *graspable*. There are two notable definitions for the word:

1. *able to understand or comprehend :* ***comprehensible***
2. *able to make concrete and perceptible; to exist in physical form :* ***embodied***

At a first glance, the scientific visualization and data physicalization communities presently take very different approaches to graspability; scientific visualizations are frequently focused on enhancing the *comprehensibility* of visualizations to enable domain scientists to discover new insights about their multivariate time-varying 3D spatial data, whereas most data physicalizations tend to focus on making the data appropriate for presentation and communication or even enjoying the data in a casual setting [Pousman et al., 2007]. Looking closer, we see that these definitions are not mutually exclusive; in fact, in some cases data physicalizations are *more* comprehensible than their digital counterparts while retaining the embodied advantages of physical, real-world objects [Jansen et al., 2013, Danyluk et al., 2020].

Henceforth in this dissertation, I define *graspable data* to be data that are made both more **comprehensible** and more **embodied** through physical rendering processes.

Physical rendering processes make data more **comprehensible** by leveraging the multisensory perceptual benefits of physical objects. With physical data objects, it's possible for data stakeholders to take advantage of both visual and haptic senses to actively perceive the data's shape by touching, holding, and running multiple fingers over the data, not to mention the additional material properties (such as squishiness or texture) that can encode additional data variables. Beyond the benefits of direct-touch interaction, physical data objects also enable us to make judgments of size and distance using our own physical reference frame (e.g., marking points with fingers or using a hand as a measuring tool) [Jansen et al., 2013]. Due to their existence in the physical world, physical data objects additionally enable true depth perception, and are more accessible to low-vision and blind users than digital displays [Jansen et al., 2015].

Physical rendering processes also make data more **embodied** by leveraging our lived experience in the physical world to make data more relatable. We, as humans, are intimately familiar with the workings of the physical world; physical data objects provide embodied interactions and affordances that we already understand, as opposed to learning new digital user interfaces to interact with the data using input and output devices. Additionally, physical data objects offer cognitive benefits for learning and understanding. In one view of embodied cognition, off-loading cognitive work into the environment is a common strategy that can be applied to data analysis with physical objects, for example, using body movements to minimize the amount of memory required [Wilson, 2002]. Ultimately, physicalizations have the potential to make data perceivable and engage audiences in ways digital data displays alone cannot, through their embodied form factor [Jansen et al., 2015].

Consider the extended reality visualization in Figure 1.2. This visualization enables collaborators to *comprehend* a multivariate time-varying 3D spatial dataset by using *embodied* interactions to make sense of ocean current flows and chlorophyll densities in relation to the depth of the ocean and other variables. The physical object encoding ocean depth data provides a natural interaction to touch points of interest, for example to mark depth at one point and compare with another. It also gives a physical reference frame as an object of common discussion amongst collaborators, enabling multiple individuals to refer to the same location on the physical data object with certainty. Additionally, the augmented reality visualization affords 3D depth cues much like the real world, enabling collaborators to look around the visualization, and get closer or farther away as they would with any physical object.

While a handful of individual examples exist of visualizations in the style of Figure 1.2 that combine both dimensions of data graspability (e.g., Tangible Landscape [Millar et al., 2018], Planwell [Nittala et al., 2015]), scientific visualization and data physicalization as a whole have yet to take a methodical approach to advancing visualizations that are both comprehensible *and* embodied.

Figure 1.2: Collaborators examine an extended reality visualization of biogeochemistry data in the Gulf of Mexico showing nitrate and chlorophyll densities as well as ocean currents. The ocean depth data is encoded with a data physicalization and the current streamlines and density glyphs are encoded as a digital immersive augmented reality overlay.

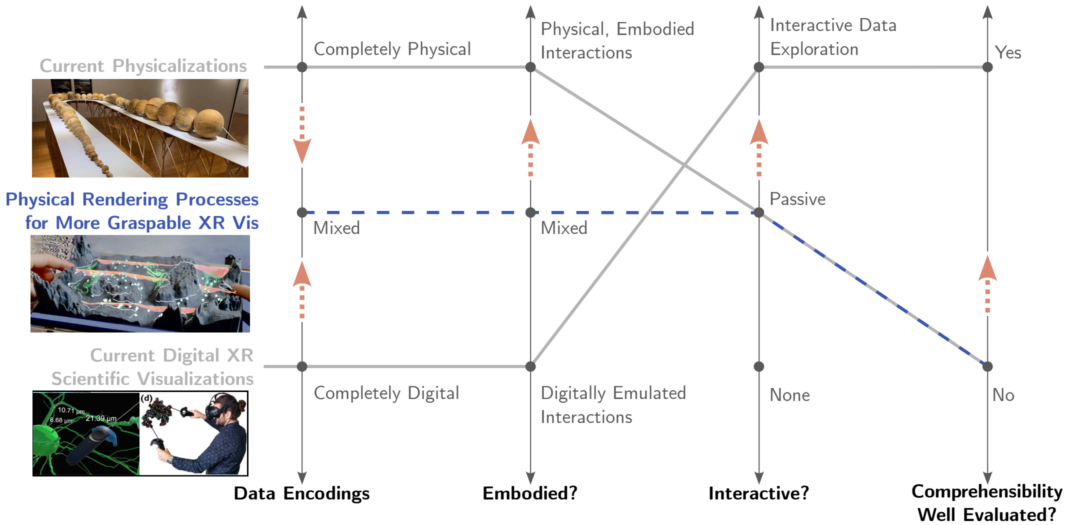

Figure 1.3: This dissertation seeks to make XR visualizations more graspable in four facets: *data encodings*, *embodiment*, *interactivity*, and *evaluation of comprehensibility*. Dotted arrows indicate the contribution of the dissertation along each facet, and the dashed line indicates the state of research prior to this work. I address the limitations of current data physicalizations (top) and current XR scientific visualizations (bottom) by introducing physical rendering processes for more graspable XR visualization (middle).

1.3 Conceptual Framework and Thesis Statement

This dissertation explores new approaches to making XR visualizations more graspable by using physical rendering processes. The conceptual framework encapsulates four facets of this work: *data encodings*, *embodiment*, *interactivity*, and *evaluation of comprehensibility*. The contributions in this dissertation advance each facet along the orange arrows shown in Figure 1.3[3], towards mixing digital and physical *data encodings*, supporting *physical and embodied interactions*, enabling *interactive data exploration*, and better understanding the *comprehensibility* of data physicalizations through *evaluation*.

This work advances the *data encodings* facet by addressing the limited physicality of current XR scientific visualizations, and addressing the limited interactivity of current, typically static data physicalizations. The work includes multiple examples in the style of Figure 1.2 that use physical rendering processes to blur the lines between physical data encodings and digital data encodings: adding digital interactivity and extended reality graphics to data physicalizations, and using physical elements from the real world to encode digital data.

This work advances the *embodiment* facet by enabling physical, embodied interac-

[3] *Bottom image from [El Beheiry et al., 2019], top image from [Madsen, 2018] and © the respective authors. Middle images from the author's work, © 2021.*

tions with data. While current XR scientific visualizations (Figure 1.3 bottom) aim to emulate the physical interactions afforded by objects in the real world; even the best such examples lack the ability for users to run multiple fingers over a data object to touch, mark points, and use our human embodied reference frame to make judgments of size and distance. The dissertation includes examples that are created with physical rendering processes, ranging from physical interactions with scalar fields superimposed on a 3D surface to physical multi-touch gestures that interact directly with 3D data.

This work advances the *interactivity* facet by enabling interactive exploration on spatial data physicalizations. While current data physicalizations (Figure 1.3 top) naturally support passive interactivity like moving the data around, yet these physicalizations lack the interactive exploration tools like querying and filtering necessary for today's scientific sensemaking processes. The dissertation includes new approaches to sensing and reimagining classic exploration-oriented scientific visualization queries in the context of data physicalizations created with physical rendering processes.

Finally, this work advances the *evaluation of comprehensibility* facet by introducing a new evaluation for comparing user performance on spatial data physicalizations. There are many current examples of data physicalization in the literature, but for all their hypothesized benefits, there have been few evaluations that substantiate the purported benefits. Current XR scientific visualization systems and techniques are comparatively well-evaluated, with *time*, *errors*, and *confidence* being frequently used metrics to compare user and algorithm performance between techniques. This dissertation includes an evaluation comparing performance metrics of users completing spatial data analysis tasks on data physicalizations, with VR visualizations, and with 2D visualizations. Since this is the first such study, but we note that there's still much work to be done before the comprehensibility of data physicalizations is evaluated to a level matching XR visualizations. Nevertheless, we advance this facet to better understand the comprehensibility of data physicalizations.

Altogether, in this dissertation I present new approaches to each of these facets that use *physical rendering processes* to overcome the limitations of current data physicalizations and extended reality visualizations, and to make data *more graspable* – both more comprehensible and more embodied.

> *New approaches to extended reality visualization that encode data using physical rendering processes enable better comprehension of and more embodied relationships with spatial data by leveraging the perceptual benefits and tactile interactions that are missing from current digital visualizations.*

1.4 Key Contributions of this Dissertation

In this dissertation, I present the following contributions for making data more graspable using extended reality visualization:

- A systematic design exploration of the first 3D printed visualizations encoding scalar data as glyphs on a 3D surface.
- A software architecture and user interface enabling the design of multivariate 3D spatial data visualizations using handcrafted physical media.
- A new approach to querying 3D spatial data using multi-touch input directly on a data physicalization.
- An empirical study on the effectiveness of data physicalizations for spatial data analysis tasks when compared with state-of-the-art virtual reality and 2D geographic information science visualizations.

1.5 Overview of the Dissertation

This dissertation is organized as shown in Table 1.1. Chapter 2 presents related work on making data more graspable with physical rendering processes. Chapter 3 presents a systematic design exploration of a new type of data physicalization, using physical rendering process to transform digital data into a data physicalization. Chapter 4 presents an inverse approach to the physical rendering process where physical artifacts from an artist's studio are used as visual encodings in a digital data visualization. Chapter 5 presents a new approach to making data physicalizations interactive by tightly integrating digital and physical visualizations with extended reality. Chapter 6 presents an empirical evaluation comparing user performance between visualization modalities: 2D, VR, and data physicalization. Finally, Chapter 7 concludes the dissertation with discussions on the lessons learned from the chapters, and the future of physical rendering processes for making extended reality data visualizations more graspable.

The work contained in Chapters 3-6 of this dissertation is the result of multiple projects among various groups of multidisciplinary collaborators, mentors, and students. I acknowledge these individuals and collaborations at the beginning of each chapter via footnotes, and I provide references to each work published with my co-authors.

Ch.	Data Source	Phys Rend Process	Data Presentation Modality
3	Generated, representative of medical data	Forward	Data Physicalization
4	Physical phenomena via remote sensing, supercomputer simulation	Inverse	Digital VR visualization
5	Physical phenomena via remote sensing, supercomputer simulation	Both	Digital immersive AR visualization, data physicalization
6	Physical phenomena via remote sensing	Forward	Physicalization, 2D desktop, Fishtank VR

Table 1.1: Chapters 3-6 are organized by their usage of physical rendering processes.

chapter two

Related Work

2.1 Making Data More Comprehensible **14**
2.1.1 Visualization Design 14
2.1.2 Extended Reality Visualizations 17
2.2 Tangible and Embodied Interaction **18**
2.3 Physical Rendering Processes **19**
2.3.1 Forward: Digital to Physical 20
2.3.2 Inverse: Physical to Digital 22

2.1 Making Data More Comprehensible

In Chapter 1, we encountered multiple approaches to making data more graspable through physical rendering processes. In this chapter, I provide related work on the overarching topic of making data more graspable, while related work specifically related to each chapter may be found within that chapter.

One of the ultimate goals of research in visualization is to make data more comprehensible. That is, visualization "*provide[s] visual representations of datasets designed to help people carry out tasks more effectively.*" [Munzner, 2014]. In this section, we review different approaches to making data more graspable, including designing visualizations to be more comprehensible, and changing the visualization modality to better suit the data; for example, using immersive digital visualizations for 3D spatial data.

2.1.1 Visualization Design

The design of visualizations is essential to their utility for making data more comprehensible. Since the beginning of visualization research, many guidelines have been developed for designing effective visualizations, including use of color, texture, symbols and glyphs, animation, and 2D vs. 3D modalities, just to name a few. Until recent decades, research in to the design of visualizations was largely focused on novel proofs of concept, and few rigorous scientific evaluations of visualization techniques and software existed prior to the year 2000 [Kosara et al., 2003]. Today, the basis of visualization research has shifted to empirical evaluation techniques, and a majority of current visualization design recommendations come from user studies about how humans perceive data. Taking a different approach to visualization design, artists have been involved in visualization as well: from Leonardo da Vinci's drawings visualizing concepts far ahead of their time [Kaplan, 1997] to present-day conferences that share the latest works combining art and visualization (e.g., the IEEE VIS Arts Program), "renaissance teams" of artists, scientists, and technologists collaborating to make visualizations have had extensive success stories throughout their history [Cox, 1991]. There are too many facets of designing visualizations to fully enumerate here, so in this section we focus the discussion on the design aspects most relevant to the Chapters in this dissertation, including the use of color and glyphs. We discuss color and glyphs from both a perceptual science standpoint (making visualizations more comprehensible) and an artistic standpoint (making visualizations more relatable), and we note that considerations brought up by both of these viewpoints are rarely mutually exclusive; in fact the visualization design decisions made by artists' expertise are often validated by perceptual science [Samsel et al., 2015].

Color

Color is one of the comprehensible and most widely used techniques to encode variables in both spatial and non-spatial data, and thus it has received a correspondingly large amount of attention and evaluation in the literature. Frequently, data are encoded via colormaps that vary luminance, saturation, and/or hue to represent specific data values, and a majority of research on color in visualization is devoted to devising "optimal" colormaps for particular datasets. While it is unlikely that a colormap exists that works equally well for all datasets [Samsel et al., 2019b], numerous user study experiments have been conducted to elucidate various properties about color in individual contexts, such as univariate colormaps [Ware, 1988] and foreground vs. background color [Wang et al., 2008]. There is also precedent for numerically optimizing colormaps based on perceptually-defined color spaces (e.g., CIELuv) [Fang et al., 2017]. Furthermore, there exist many surveys on the usage of color in visualization [Silva et al., 2011, Zhou and Hansen, 2016] which provide guidance for the design of colormaps

Complementing the perceptual science approach to color, artistic color theory has also been applied to make visualizations more comprehensible to scientists and the public. In 1988, the artist Donna Cox developed groundbreaking colormaps to encode the high-dimensional simulation results produced by a supercomputer [Cox, 1988]. Recent research by Samsel et al. has shown that certain colormaps are not only better to use for certain datasets from a pure perceptual standpoint but also provide affective and associative cues that intuitively indicate what colormaps are representing [Samsel et al., 2019b, Samsel et al., 2018a]. There have also been many guides for artists and visualization practitioners on effectively applying principles from artistic color theory to visualizations [Rhyne, 2012]. In Chapter 4, we build on the related work in artist-designed colormaps for scientific visualization.

Glyphs

Glyphs are particularly useful technique for encoding discretely sampled variables at specific points in spatial data visualizations while also enabling visual inspection at different scales [Kindlmann and Westin, 2006]. In particular, glyphs are ideal for depicting multiple variables at once, for example *pressure* and *temperature* [Lie et al., 2009]. There are various styles of glyphs, including geometric primitives like those found in general purpose visualization software [Ahrens et al., 2005, Childs et al., 2012], superquadrics [Schultz and Kindlmann, 2010], and more novel forms like flowers [van Onzenoodt et al., 2022] or faces [Morris and Ebert, 2000]. Glyphs have been used for many types of domain-specific data, but one application area that's received special interest for glyph-based visualization is diffusion-tensor medical imaging [Zhang et al., 2016, Kindlmann and

Westin, 2006, Schultz and Kindlmann, 2010]. Direction-based glyphs are ideal for this domain because of their ability to depict the individual axes of each tensor along with additional variables like the density value [Kindlmann and Westin, 2006, Farooq et al., 2016]. Glyphs ultimately work towards making complex, multivariate data more comprehensible and the existing examples so far in the literature provide a start to rigorous evaluation for glyphs, but as a past survey points out, there are still many open questions regarding the ideal glyph design for individual datasets [Fuchs et al., 2017].

Beyond the computationally-defined techniques above, artists have worked in parallel to create glyph-based visualizations as well. Keefe et al. describe a four-stage approach to leveraging artists' expertise in visual design for scientific visualizations, including rapidly sketching glyphs designs first on paper then in virtual reality to afford better comprehension of spatial, multivariate data [Keefe et al., 2008a]. Similarly, Visualization-by-Sketching enables artists to directly sketch glyphs onto a 2D spatial data visualization using a digital stylus [Schroeder and Keefe, 2016]. Work has also been done to evaluate the affect and associative relationships of different artist-designed 3D glyph styles; similar to affective color, the goal of this work was to enable the use of different glyphs for semantic and emotional meanings to make 3D spatial data visualizations [Zeller et al., 2022]. The Artifact-Based Rendering technique (ABR) enables an artist to create 3D glyphs out of physical materials in their studio (e.g., clay, pebbles, rice), transfer the glyphs into a digital format, and use them to encode data [Johnson et al., 2019b]. This groundbreaking technique for leveraging artists' training and expertise in visual design is the starting point for our work in Chapter 4.

Ultimately, scientific visualization uses color, glyphs, and many more visualization design techniques and more to make data more comprehensible. Color can be combined with texture to show fluid flow simulation results [Laramee et al., 2004b], stream surfaces and streamlines can be extracted from similar flow data [Laramee et al., 2004a], and streamlines with varied thickness can be combined with glyphs, color, and texture to make more effective weather maps with wind direction markers [Pilar and Ware, 2013]. Glyphs can be combined with streamlines and cutting planes to highlight and filter complex vector fields [Stevens et al., 2020]. Others have even established perceptual guidelines for creating volumetric visualizations of 3D spatial data based on empirical user studies [Boucheny et al., 2009]. These works are by no means a comprehensive list of the techniques and tools available to scientific visualization practitioners, but they show a bird's eye view of the techniques that are most relevant to the remainder of this dissertation.

2.1.2 Extended Reality Visualizations

Extended reality visualizations are particularly well suited to visualizing 3D spatial data like those found in many scientific disciplines today; they are able to depict these data without projecting them into 2D, which is often crucial for comprehending complicated physical phenomena that domain scientists are trying to understand [Bryson, 1996]. Virtual reality, even from its early days, has enabled scientists to view their 3D spatial data (e.g., physics data [Bryson, 1992] or fluid flow analysis [van Dam et al., 2002]) immersively by using natural head movements, and enables better data understanding through interactive querying, filtering, and moving the data around via spatial input devices like wands [Cruz-Neira et al., 1993] or gloves [Fang et al., 2015]. Today, virtual reality also enables data analysts to better understand statistical and other non-spatial data via immersive analytics [Fonnet and Prie, 2021]. Augmented and mixed reality afford even more of a real-world experience to users by showing digital content overlaid on top of the physical world; this enables spatial experiences to involve components from the real world, usually tangible props, for example paper image targets [Bach et al., 2018] or writing utensils [Bobrich and Otto, 2002].

Even though there are solutions in AR and VR that provide some level of tactile feedback – for example, haptic rendering devices like robotic styluses [Laycock and Day, 2007] – in general these only provide a low resolution glimpse into the data (i.e., only one point at a time) and require expensive, custom hardware. These robotic, actuated solutions may have a promising future to achieve real-time, high-resolution, low-latency physical display of data, but these are not yet at the point where they're feasible to use as a part of scientists' everyday sensemaking processes due to their associated specialized hardware and software costs.

As a result, the most widely adopted examples of extended reality visualization in the sciences rarely use even the simplest of these attempts at recreating the tactility privileged by physical sensemaking processes. In fact, multiple recent survey papers across various scientific disciplines (e.g., geographic information science [Çöltekin et al., 2020], biomedical science [Venkatesan et al., 2021], and archaeology [Bonis et al., 2022]) show exclusively examples that only utilize the built-in controllers that ship with commercial, off-the-shelf extended reality devices. While the standard controllers are usually the most cost-effective solution, they lack the capacity to support the type of exploration and tactile feedback that even simple everyday physical objects exhibit in the real world. Thus, we are motivated to explore extended reality visualizations that use physical rendering processes to enable embodied interactions like touching, moving and rotating, and running multiple fingers across the data to enable better comprehension of data in scientific sensemaking processes.

2.2 Tangible and Embodied Interaction

The work in this dissertation also builds on the field known as Tangible and Embodied Interaction (TEI). TEI focuses on the "*implications and new possibilities for interacting with computational objects within the physical world*" [Hornecker, 2011]. Though it is a relatively new as a research area, TEI has a rich history combining work from such fields as cognitive science, industrial design, the arts, and computer science to create tangible, embodied interactions with computers and data [Hornecker and Buur, 2006]. TEI builds on the theories of embodied cognition, in particular the view that the physical world and environment plays a crucial role in how we, as humans, think about and perceive information [Wilson, 2002], and holds many of the same principles true as data physicalization does: TEIs and physicalizations are both located in physical space, and they share this space with users, affording interactions that are experienced in relation to the user's body [Hornecker et al., 2023]. We also note that TEIs share a common motivation with extended reality visualization and immersive analytics, namely the key sense of "embodiment" that walking around the data or moving the data in relation to oneself can provide [Büschel et al., 2018].

Tangible User Interfaces (TUIs) are an example of a technology that enables Tangible and Embodied Interactions. Early pioneering works from the literature include Bricks [Fitzmaurice et al., 1995] and Tangible Bits [Ishii and Ullmer, 1997], which both show capabilities for blending the digital and physical worlds far beyond their time, using every day physical objects instrumented with trackers to control digital content. More recent examples from the literature include everything from combining TUIs with virtual reality to make the VR experience more embodied [Chang et al., 2017] to new teaching tools to demystify machine learning algorithms [Kaspersen et al., 2021]. These examples all leverage users' relationships with physical objects in the real world to make interactions with digital content more embodied and relatable [Shaer and Hornecker, 2010].

Inspired by successful examples of TEIs from non-visualization contexts, the data visualization research community has also adopted principles from TEI to interact with data via physical props. Usually, these props are an abstract representation of a 3D cursor that can make queries or annotations in space, for example using a physical pen to make AR annotations on a map [Bobrich and Otto, 2002], filtering 3D scatter plots in AR with physical cutting planes [Bach et al., 2018], manipulating and querying geophysical data with a physical globe [Satriadi et al., 2022], and using the affordances of a physical piece of paper to zoom, filter, and add details on demand to AR visualizations [Tong et al., 2022]. In these typical usages of TEIs for visualization, the tangible objects enable a new level of physicality to otherwise entirely digital data displays, but the focus of the visualization remains on the digital data, and the TEIs rarely encode data themselves.

Data physicalization takes TEIs with data to a new level, enabling users to explore the data through movement in relation to their own embodied reference frame, taking advantage of the ability to measure distance with body parts, to change perspective by physically moving the object or self, and to walk through a visualization [Hornecker et al., 2023]. This enables users to take advantage of active perception while exploring the data, which is beneficial to comprehending data due to visual perception being highly integrated with movement and haptic perception [Gibson, 2014]. Additionally, data physicalizations enable users to perceive the data with unambiguous stereoscopic depth cues (i.e., the real world), as opposed to the simulated depth cues provided by digital extended reality displays, which can be incorrect for many reasons, including the accommodation-vergence conflict [Kramida, 2016] and incorrect inter-pupilary distance estimations [Best, 1996]. Beyond their perceptual and cognitive advantages, data physicalizations also enable many implicit advantages by their use of physical objects: improved memorability [Lederman and Klatzky, 2009], emotional connection with the data [Wang et al., 2019], social connection with other individuals discussing the data [Hogan and Hornecker, 2012], and rich metaphors enabled by their material properties [Hornecker et al., 2023]. In this dissertation, I aim to take advantage of the prior work on both data physicalizations and Tangible and Embodied Interactions as a whole, looking at these works through the lens of creating more comprehensible and embodied visualizations with physical rendering processes.

2.3 Physical Rendering Processes

In this section, we discuss the related work on bringing data into the real world through the lens of the expanded definition of the physical rendering process introduced in Chapter 1. To begin, we contrast this definition with related concepts from extended reality and computer graphics literature.

The process of bringing data into the physical world is distinct from *haptic rendering* originating in the extended reality literature [Salisbury et al., 2004] and *physically-based rendering* (PBR) from computer graphics [Pharr et al., 2016]. Haptic rendering attempts to replicate the tactile feedback experienced with physical objects by using complex actuation mechanisms like robotic gloves [Fang et al., 2015] or force-feedback styluses [Jackson et al., 2012a], and while some of these technologies enable a passable approximation of touching a physical object, they are not at a point where they support rich touch interactions like those enabled by physical objects, such as holding, rotating, marking multiple points, and using our embodied reference frame to make size and distance judgments. Physically-based rendering uses the mathematical expression of laws of physics to generate realistic computer graphics images, and while these techniques have produced stunning digital images that can fool the viewer into believing they were taken in the real world [Ulbricht et al., 2006], the focus of PBR is specifically on the creation of *digital* images and does not encompass making any images or data *physical*.

With these distinctions in mind, we introduce related work on the "forward" physical rendering process of making data physicalizations through digital fabrication technologies, as well as the "inverse" physical rendering process of bringing elements from the physical world to life in digital visualizations.

2.3.1 Forward: Digital to Physical

Data physicalizations are most commonly a "one-shot" through the physical rendering process. That is, these static, passive physicalizations take one instance of one dataset at one moment in time and bring it into the physical world through digital fabrication [Jansen and Dragicevic, 2013]. Examples in this style include molecular data physicalizations [Bailey et al., 1998], data furniture [Duffy, 2014, Segal, 2011], and sculptures depicting flight data [Spitz, 2013] or physical activity data [Stusak et al., 2014]. While technically only "rendered" once, other physicalizations incorporate interactions with the environment they're situated in, such as wind sculptures [Kahn, 1996] and living physicalizations using plants [Meier, 2017, Sigitas Guzauskas, 2018]. Lastly, some physicalizations feature manual physical interactivity such as rearrangeable parts; for example Vologram enables data stakeholders to view an entire volumetric visualization or individually view slices along three axes by pulling them out of the visualization [Pahr, 2021]. Our work in Chapters 3 and 6 fall into the *static* or *passive* category of physicalization.

In contrast to these passive physicalizations, *active* physicalizations have been introduced to enable user interactivity and animation capabilities like those frequently necessitated by today's purely digital visualizations. In the early 2000s, a prototype active physicalization that featured moving, actuated vertical columns augmented was used for depicting solid terrain models; this apparatus was augmented with a digital projection that enabled further data to be displayed on the dynamic physical surface [Schmitz, 2004]. Similar active, robotically actuated approaches have been used since then for similar terrain visualizations [Follmer et al., 2013] and analytics of statistical data [Taher et al., 2016]. Other types of active physicalizations also exist, including the use of wheeled micro-robots that position themselves into visualizations on a tabletop [Le Goc et al., 2019] and acoustic levitation of physical objects to form 2D or 3D visualizations [Omirou et al., 2016, Gao et al., 2023]. These impressive active physicalizations are changing the way we physicalize data, but they presently are quite expensive and difficult to create, calibrate, and maintain due to the sheer number of moving parts involved (for example, the EMERGE display used in [Taher et al., 2016] uses a 10x10 grid of actuated columnar "pixels" that draw a significant amount of current and are prone to overheating). In this dissertation we turn to other, more cost-effective approaches that are presently easier to fabricate and maintain.

An alternative to active physicalizations that still enables interactivity and dynamic content on physicalizations is augmented physicalizations. While many active physicalizations are also augmented with additional digital content (e.g., projection [Follmer et al., 2013] or light [Taher et al., 2016]), it is more common for statically fabricated physicalizations to be digitally augmented. Such is the case for Gillet et al.'s early molecular model visualization that uses a webcam video feed to augment a physicalization with additional, digital dynamic data [Gillet et al., 2005]. In the literature, projection has been the most common means of augmenting a physicalization, likely due to its innate capacity to support shared viewership that other technologies don't have (e.g., immersive AR) [Djavaherpour et al., 2021]. Some examples of projection-augmented physicalizations include TanGeoMS [Tateosian et al., 2010] and PARM [Priestnall et al., 2012], which have been used for landscape modelling, as well as a visualization depicting energy usage [Kirshenbaum et al., 2020]. Mobile augmented reality can also be used for visualization; for example, augmenting ring physicalizations with statistical data [Bilgili, 2017] and touch-enabled glyph-based spatial visualizations [Chen et al., 2020]. With the development and improvement over time of immersive AR devices, examples of immersive AR augmentation are also becoming more prevalent, for example PLANWELL [Nittala et al., 2015] uses a tracked physical model and immersive AR visualization to facilitate more effective communication between teams for petroleum well drilling projects. In our work, Chapter 5 uses immersive AR to augment a touch sensitive data physicalization.

We also note that some definitions of *data physicalization* are more restrictive than others. For example, some definitions would not consider the physical digital elevation data like those presented in Chapter 6 *data physicalizations*; rather, it could be considered a *physical model* [Dragicevic et al., 2021, Jansen et al., 2015]. We argue that even physical objects representing "simple," single-variable spatial data like digital elevation data would be considered a data physicalization; after all, it uses the *geometric and material properties of a physical artifact to encode data* [Jansen et al., 2015]. Thus, we adopt a less restrictive definition for the rest of the dissertation: any physical object that encodes data of any kind with its geometry or material properties is considered a data physicalization.

2.3.2 Inverse: Physical to Digital

In addition to the "forward" physical rendering process of bringing digital data into the physical world via data physicalization, this dissertation also explores the "inverse" physical rendering process where elements from the physical world are transformed into the digital domain and used as encodings in digital data visualizations. Inspired by artistic work created with physical media like that of Nathalie Miebach [Miebach et al., 2022], the inverse physical rendering process aspires to enable a richer visual vocabulary in data visualization by meeting artists where they're frequently the most comfortable working: in the physical world. Work in the style of Scientific Sketching [Keefe et al., 2008a] and Artifact-Based Rendering [Johnson et al., 2019b] has made it possible for artists in their studio to quickly create hundreds of design iterations on the design of even a single data variable with physical media they work with on a daily basis. This not only expands the visual vocabulary for scientific visualization but also – through informed use of color, line, texture, and form – has the potential to enable scientists to comprehend more of their data, more effectively, and in a more embodied way [Johnson et al., 2019b]. Chapters 4 and 5 elaborate on our notion of the inverse physical rendering process. Specifically, Chapter 4 explains the process through which it is possible for artists to begin their design process in the physical world, with physical media, and end with a multivariate 3D virtual reality scientific visualization, and Chapter 5 addresses the completion of the cycle: rendering the artist-created physical artifacts back into the physical world using immersive augmented reality.

Though not directly related to the work in this dissertation, Dietmar Offenhuber's concept of *autographic visualization* [Offenhuber, 2020] also fits into the inverse physical rendering process. Autographic visualization takes physical phenomena, instead of data, as the starting point in the visualization process. Whereas traditional data visualizations take data recorded about a physical phenomenon through sensors, transform it digitally, and finally re-create visual representations of the phenomenon in question, autographic visualizations instead seek to utilize the physical phenomenon as data about itself. Autographic visualization is still an emerging field, and though autographic visualizations rarely end up in the digital domain, their applicability to our work is clear for their priority of using physical processes to develop better comprehension about physical processes.

In summary, related work on data visualization has strived to make data more comprehensible through improvements to the design of visualizations through both empirical means (user studies) and artistic means (artist training). And, through physical rendering processes, we can make data more graspable in extended reality visualizations. For the remaining chapters in this dissertation, we expand on the forward and inverse physical rendering process for making data more graspable.

chapter
three

Designing Glyph-Based Spatial Data Physicalizations

3.1 Introduction 24
3.2 Methodology 24
3.2.1 Research Through Design 25
3.2.2 Current Visualization Practice as a Starting Point 25
3.3 Designing Glyph-Based Spatial Data Physicalizations 27
3.3.1 Visual Channel 28
3.3.2 Glyph Shape 29
3.3.3 Glyph Size and Data Remapping 29
3.3.4 Glyph Alignment With Surface 30
3.3.5 Glyph Placement On Surface 30
3.4 Fabricating Glyph-Based Spatial Data Physicalizations 30
3.5 Results and Discussion 32
3.6 Conclusion 33

3.1 Introduction

[1]So far in this dissertation, we've seen the important role that the science of data visualization can play through collaborations with expert domain scientists to devise more effective sensemaking processes. The work in this chapter is motivated by one such collaboration between our visualization lab and a team of medical device engineers studying the effects of heart implants on blood flow.

Frequently, medical device engineers like those we collaborate with need to visualize variables *in conjunction with* the 3D spatial geometry; for example, they might choose to visualize the blood pressure (scalar variable) overlaid on the surface of the right atrium of the heart. Visualizing these data is especially challenging because of the 3D spatial and multivariate nature of the data; typically 2D representations shown on a desktop monitor with simple keyboard and mouse interactions are not sufficient for the engineers to sufficiently answer their driving scientific questions [Nielson et al., 1997]. Some medical device engineers have turned to extended reality visualizations like virtual and augmented reality [Venkatesan et al., 2021], while some even use advanced XR technologies like haptic feedback [Jackson et al., 2012a] or physical props [Jackson et al., 2013]. Yet, despite these technological advancements and new visualization techniques, engineers still struggle to answer the questions they have of their data, which require an intimate understanding of the 3D surface geometry *and* the data variables on top of the surface.

We interpret the challenges visualization techniques have visualizing these data as opportunities for data physicalization. In contrast to even the best extended reality visual and haptic displays, data physicalizations provide zero-latency interactions with data that are only limited in resolution by the fabrication device used; they also engage the sense of touch and provide true stereoscopic depth cues. In this chapter, we present a systematic design exploration of the first data physicalizations that combine 3D surfaces with glyphs encoding a scalar field on the surface. This exploration is motivated by the visualization needs identified by our collaborators, and builds on related work from the scientific visualization, information visualization, and data physicalization communities.

3.2 Methodology

Our research explores the design space for spatial data physicalizations that overlay scalar fields on 3D surface geometries and combines practices from computer science

[1] This chapter is based on work published as a chapter in the *Making Data* book [Gwilt, 2022], under the title *Hybrid Data Constructs: Interacting with Biomedical Data in Augmented Spaces* [Keefe et al., 2022b], as well as a juried pictorial paper presented at the IEEE VIS 2018 workshop on data physicalization [Herman and Keefe, 2018].

research and design methodology. In particular, we begin with a process inspired by the "Research through Design" framework [Zimmerman and Forlizzi, 2014], which we use to guide our exploration and filtering of current visualization best practices for scalar data and 3D surfaces.

3.2.1 Research Through Design

Research through Design (RtD) is a relatively recent reinterpretation of design practices as applied to scientific research, specifically in Human-Computer Interaction; RtD "*draws on design's strength as a reflective practice of continually reinterpreting and reframing a problematic situation through a process of making and critiquing artifacts that function as proposed solutions*" [Zimmerman and Forlizzi, 2014, p. 167]. This chapter is focused on a *process of making and critiquing* data physicalizations that advance the state of visualization research and help address the needs of our collaborating medical device engineers.

In the suggested RtD fashion, we begin the process by *selecting* the problem: *physically render scalar fields in conjunction with 3D surfaces*. Then, based on literature in scientific visualization and data physicalization, we *design* several possibilities for visualizing scalar fields on 3D surfaces. After fabricating a series of the design possibilities, we *evaluate* the designs based on considerations from the literature and lab members' expertise with scientific visualization. Lastly, we *reflect* on the effectiveness of each design, and *repeat* the process with new physicalization designs.

3.2.2 Current Visualization Practice as a Starting Point

Our methodology begins with current visualization best practices. To inform the early designs for data physicalizations that depict a 3D surface geometry and a scalar data field, we began by identifying the axes of the design space based on prior literature in the visualization community. We explored many options to visualize the scalar data overlaid on the surface, including colormaps [Samsel et al., 2019b, Zhou and Hansen, 2016], glyphs [Ropinski et al., 2011, Ang et al., 2019], and even surface texture [Kim et al., 2003, Johnson et al., 2019b].

We used Munzner's consolidated effectiveness of visual channels [Munzner, 2014, p. 102] as a starting point to show the design choices available for overlaying a scalar field on a 3D surface From this list, we narrowed down the possible options to visualize the scalar data on top of the surface. Since our use case (medical device engineering) requires data that are inherently 3D and spatial, channels like position and depth are unavailable to use. Due to limitations of the commercial off-the-shelf FDM printers used during prototyping, other channels like color luminance and saturation are not available

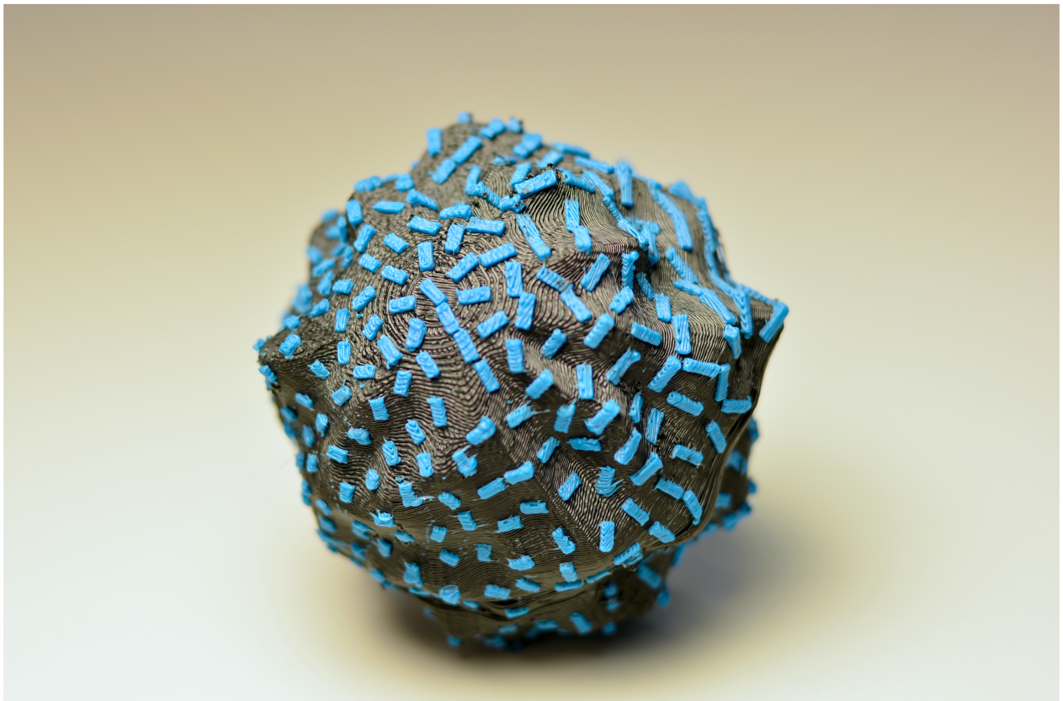

Figure 3.1: Current design for glyph-based spatial data physicalizations

either. Reducing the list to a bare minimum of feasible options for 3D spatial data and the digital manufacturing technologies available, we ended up with the *length*, *angle*, *area*, *curvature* channels. The channels used are discussed further in Section 3.3.

Based on the visual channels available, we determined that a glyph-based visualization was a good fit for these scalar data, since glyphs can have *length*, *angle*, *area*, *curvature* and properties. Glyph-based visualization is common in scientific visualization. Typically, glyphs are used to depict multivariate or high-dimensional variables at discrete, sampled points in a 2D or 3D visualization of spatial data. 3D glyphs have taken forms like textured decals on a surface for depicting a scalar variable with a vector variable [Rocha et al., 2017], superquadric geometric glyphs for depicting multiple scalar variables simultaneously [Kindlmann and Westin, 2006], and other 3D objects overlaid on a surface for visualizing multimodal volume data [Ropinski et al., 2007]. Based on these options and other considerations like ease of interpretation [Lie et al., 2009], we identify simple geometric glyphs placed on the 3D physical surface as a starting design for physically rendering scalar fields on 3D surfaces. The remainder of the sections in this chapter focus on this starting design: *glyph-based spatial data physicalizations*.

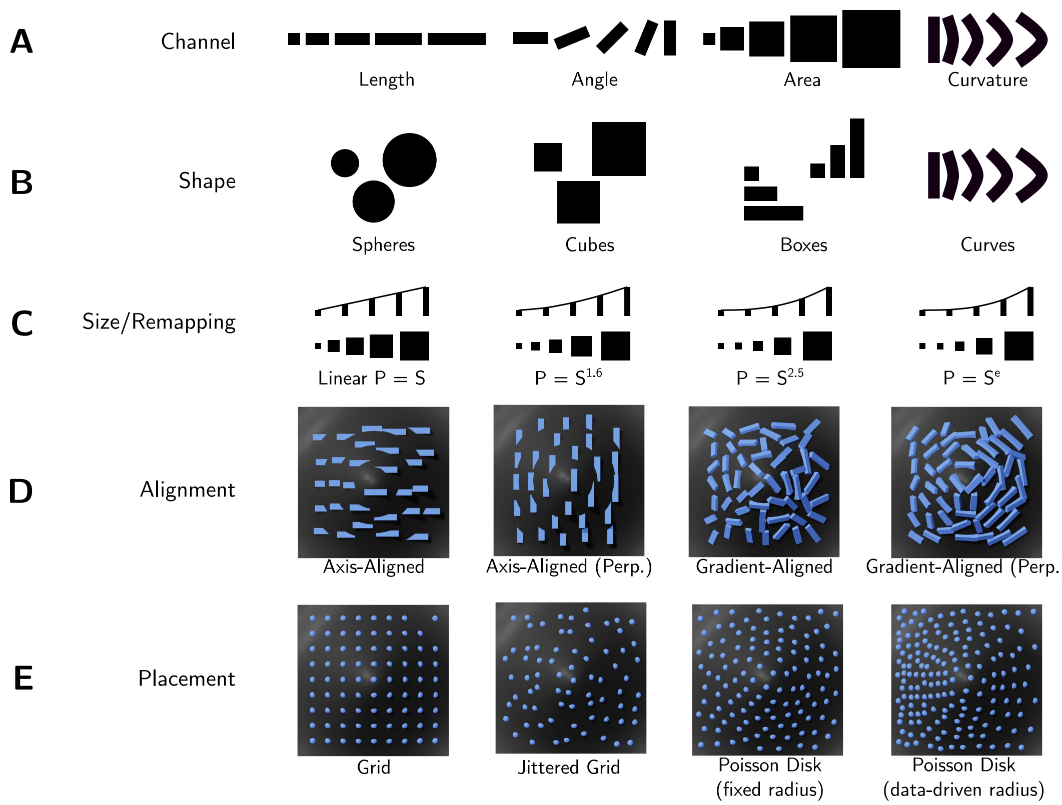

Figure 3.2: The design space for glyph-based data physicalizations of scalar values on 3D surfaces includes: the channel used to convey the scalar variable (A), the glyph shape (B), the absolute size and any size remappings applied (C), the alignment of the glyphs relative to the surface (D), and the glyph distribution/placement on the surface (E).

3.3 Designing Glyph-Based Spatial Data Physicalizations

Given the design considerations outlined so far from the literature, we created a design space for glyph-based spatial data physicalizations. The design space, shown in Figure 3.2 includes five axes: the *channel* used to convey the scalar value (from Munzner's visual channels [Munzner, 2014]), and four axes inspired by the glyph properties outlined in previous literature [Ropinski et al., 2011, Lie et al., 2009]: *glyph shape*, *glyph size*, *glyph alignment on surface*, and *glyph distribution/placement*. The remainder of this section explains the design decisions made to arrive at the final glyph-based data physicalization shown in Figure 3.1.

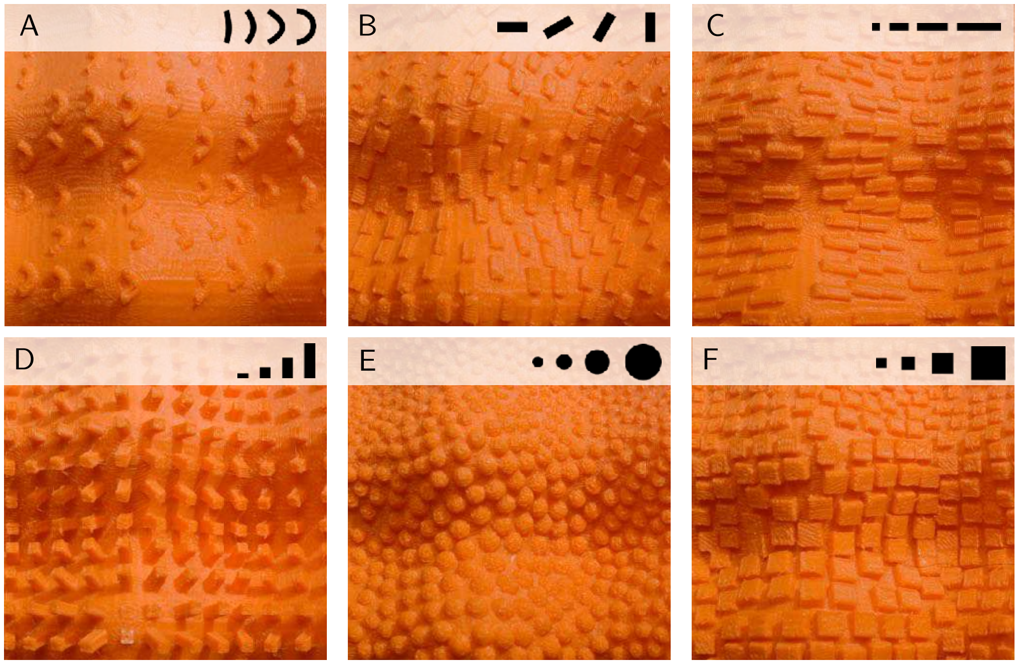

Figure 3.3: Glyph shapes used in the design process and their corresponding visual channels include: curves (curvature, A); boxes (angle, B); boxes (length, C); boxes (height, D), spheres (area, E), and cubes (area, F).

3.3.1 Visual Channel

Based on the visual channels in Figure 3.2A, we created glyph-based data physicalizations that used each of these channels to depict the scalar data field.

Initially, the *area* channel was used (Figure 3.3D and E). Area-based symmetric glyphs are widely used in scientific visualization because they are well-suited to 3D spatial data; they convey both the data value(s) and the point in space. However, they make it difficult to interpret the true values of data due to the nonlinear perception of size as a variable [Lie et al., 2009, Ropinski et al., 2011, Jansen and Hornbaek, 2016], so ultimately we sought other channels.

Angle and *curvature* are also valid channels (Figure 3.3A and B) to encode data. These channels are less common for depicting scalar fields scientific visualization because they seem ill-adapted to spatial data, as they tend to convey directional information, as opposed to the scalar value we are intending to show.

Lastly, the *length* channel tops Munzner's chart on the perceptibility of visual channels [Munzner, 2014, p. 102]. Glyph length is often used in scientific visualization to depict scalar data, and even parts of tensor data (e.g., medical imaging data [Kindlmann and Westin, 2006]). Figure 3.1 shows the current physicalization design using *length*-based glyphs depicting a scalar variable on a 3D surface.

3.3.2 Glyph Shape

The choice of glyph shape, shown in Figure 3.2B, depends heavily on the choice of visual channel: notice in Figure 3.3C how the box-shaped glyphs are not able to represent any curvature and still maintain their box shape, and spheres are not able to be used for length, angle, or curvature. During design iterations, further concerns arose for the use of spheres: some individuals seem to experience discomfort when viewing the physicalizations that used spheres. This may have been triggered by trypophobia, the fear of holes that can happen, for example, when looking at a honeycomb [Le et al., 2015]. In our design iterations, the box-shaped glyphs did not induce such discomfort on anyone who used them. Besides general user comfort, box-shaped glyphs are also advantageous because they are able to reflect properties of the surface below them instead of occluding them; for example, they can show the surface normal and gradient. Ultimately, the current glyph design shown in Figure 3.1 uses box-shaped glyphs.

3.3.3 Glyph Size and Data Remapping

Before mapping any scalar variables to the size of glyphs, two important Prior work has established that certain absolute glyph sizes are best to perceive; for 2D scatter plots considerations need to be made: the absolute size of the glyphs, and any remappings that should be applied to the data. For these considerations, we again turn to established practices in the visualization community. Prior work has established that certain absolute glyph sizes are best to perceive; for 2D scatter plots glyphs should have visual angles between 0.072°and 0.573° [Li et al., 2010]. Interpreting these guidelines at a comfortable examination distance of 25cm, our current glyph design ranges between radius 0.314mm and 2.5mm.

Absolute radius being resolved, we now turn to the data mapping, shown in Figure 3.2C. It has also been established in the literature that symbol size is not perceived linearly. In fact, for scatter plots with circular symbols, size judgments have been found to follow a power transform of $P = S^{2.5}$, where P is the value perceived by the user and S is the original size value (radius of the symbol) [Li et al., 2010]. The size of spherical physical data representations has been found to follow a power law of $P = S^{1.6}$, whereas physical bars can be treated with linear perception $P = S$ [Jansen and Hornbaek, 2016]. The series of 3D prints showing comparisons of different sizes and power mappings can be seen in Figure 3.2C, and the current design using the length of physical box-shaped glyphs is shown in Figure 3.1.

3.3.4 Glyph Alignment With Surface

The alignment of glyphs on the surface is also dependent on the type of glyphs used. For example, spheres look the same regardless of the orientation they use on the surface, whereas the alignment of cubes, boxes, and even more asymmetric shapes matters significantly. Figure 3.2D identifies the three options for aligning 3D glyphs on a surface. Axis-aligned glyphs are simple to implement, but the results show inconsistent parts of glyphs other than spheres, and they don't allow the glyph to reflect any information about the underlying surface, which is important to the data stakeholders. Normal-aligned glyphs reflect some additional information about the surface they are placed on, and they give a true "outward" facing direction for any glyph shape including irregular glyphs, but only specifying one basis vector for rotation leaves the glyph orientation underconstrained. Thus, we introduce a second vector: the surface gradient at a given point. In our current design, we use a technique inspired by Kim et al. [Kim et al., 2003] to calculate the surface gradient in the principal direction. We also note that these glyph alignments can be expanded or replaced if the data that need to be visualized are vector data instead of scalar data.

3.3.5 Glyph Placement On Surface

The last key consideration for overlaying glyphs on a 3D physical surface is the placement or distribution of the glyphs. Figure 3.2E shows four different options for glyph placement. Glyphs may be sampled from a grid, but this has been shown to introduce bias into interpreting the visualization results (e.g., if there is an important value that was missed in between the grid points) [Laidlaw et al., 2005]. Jittered grids have been used as a solution to the grid regularity problem, yet they often have similar flaws to regular grids. An alternative approach is Poisson-Disk sampling, which provides sample points which are at least some distance r apart from one another [Bridson, 2007]. Poisson-Disk sampling can also be modified to function as a data-driven placement algorithm; glyphs can be packed near each other based on their size (similar to glyph packing, as in [Kindlmann and Westin, 2006]).

3.4 Fabricating Glyph-Based Spatial Data Physicalizations

The process for fabricating the glyph-based physicalizations can be seen in Figure 3.4. Starting with Blender Python scripts[2], a visualization designer can tune the glyph design parameters including the glyph shape, glyph size and data mapping, alignment with

[2] Code can be found at https://github.com/ivlab/BoxcarPotatoes

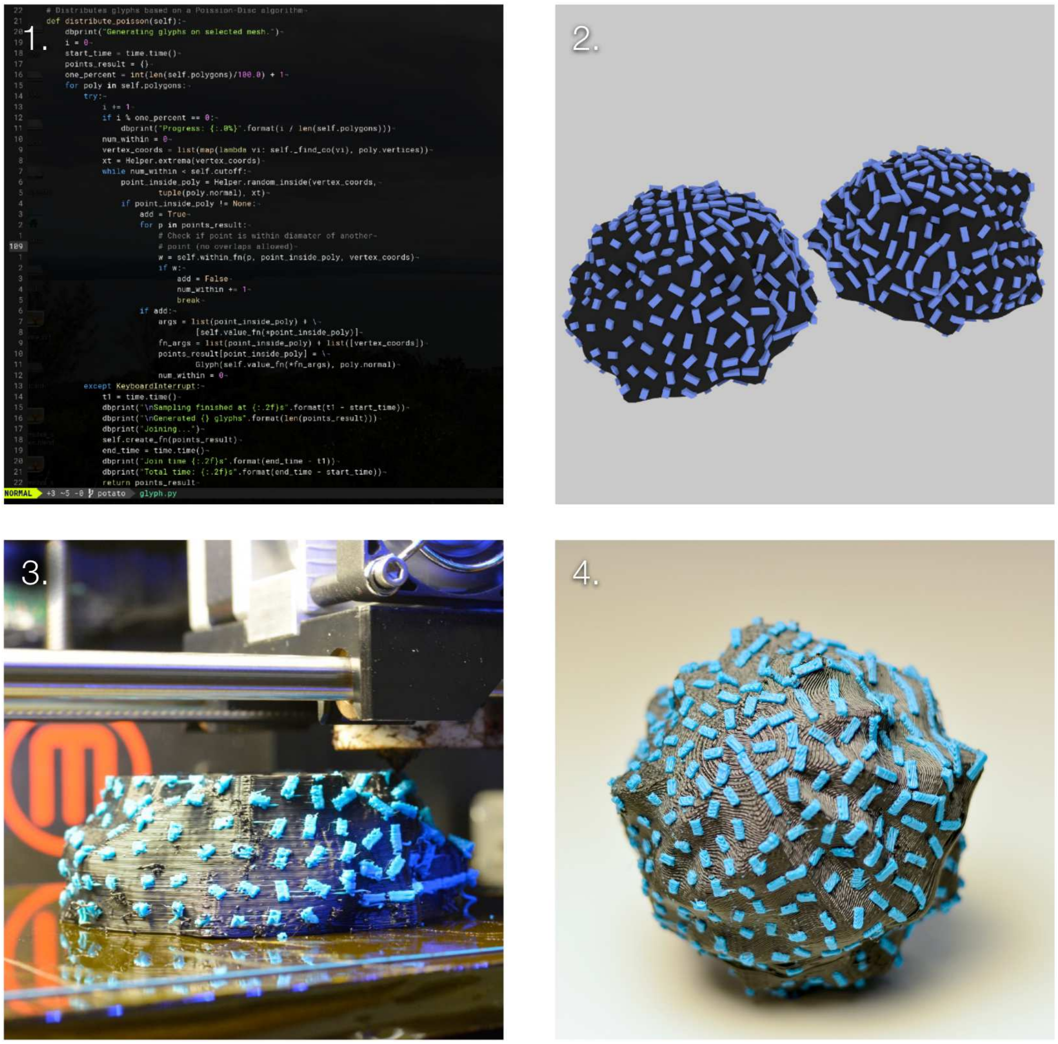

Figure 3.4: Fabrication of the designs took a four step process: (1) Blender scripts are used to generate 3D surface geometry and glyph overlays; (2) the scripts slice the potato object in half for easy 3D printing; data objects are 3D printed using a multi-color FDM 3D printer (3), the final data object after post processing and securing the two halves together (4).

surface, and placement. The scripts automatically slice any overhanging geometry; for example the current data physicalization shown in Figure 3.1 was printed in two halves and secured together after printing. After the models are generated in Blender, they can be exported to 3D printing slicer software and sent to the printer. Our fabrication process uses a dual-extruder MakerBot printer.

3.5 Results and Discussion

Though the process of evaluating glyph-based data physicalizations is ongoing and thus far we've only consulted on the designs with visualization experts from our lab, we can already report a few key findings from early lab feedback and justification from the literature.

Firstly, when visualizing a single variable with physical glyphs on top of a physical surface, it is best to use primitive shapes (i.e., spheres, boxes). In our design explorations, we found that other shapes (e.g., the "macaroni"-like curved glyphs in Figure 3.3) did not work very well to depict data and were distracting from both the scalar data they were attempting to encode and also occluded the underlying surface without providing any indications for what the surface normal or gradient was underneath each point. These findings mirror the suggestions of Lie et al. [Lie et al., 2009] regarding glyph design best practices.

Secondly, care needs to be taken when using the size of glyphs to encode a data value, especially in the physical world. In the psychology literature it is well known that area is not perceived linearly (e.g., Stevens's power law); and the same applies to digital visualizations [Li et al., 2010] and data physicalizations [Jansen and Hornbaek, 2016] as well. Ultimately in our current design, we used the length of the glyph to encode the scalar variable rather than the radius or area of the glyphs, which literature indicates is perceived nearly linearly both on digital and physical visualizations [Jansen and Hornbaek, 2016, Munzner, 2014].

Thirdly, care must also be taken when choosing a sampling method to place glyphs on the surface. Many sampling methods are available, like those shown in Figure 3.2E. In our testing we found that Poisson-Disk sampling with a fixed radius worked well for our use cases, but for more complicated data it may be necessary to resort to glyph packing or other placement methods [Kindlmann and Westin, 2006].

Lastly, since glyphs placed on a surface occlude part of the surface, it becomes necessary to find ways to convey information about the surface underneath on the glyphs themselves. In the style of Kim et al. [Kim et al., 2003], our current design aligns the glyphs with the surface normal at the sampled point, as well as the primary gradient direction of the surface underneath to give the viewer more information about the surface instead of only occluding it. An alternative method would be to extrude parts of the surface directly where the glyphs intersect; this would provide an exact replication of the important surface geometry below each glyph, but may introduce additional confounding factors if it is unclear where each glyph ends and the surface begins.

3.6 Conclusion

In this chapter, we methodically explored the design space for the first glyph-based spatial data physicalizations. Starting with findings from current visualization and data physicalization literature, we sought to devise a new type of physical data representation that would help address the needs of our medical device engineering collaborators. Our collaborators already use 3D printing to better understand anatomical surfaces, such as the patient-specific shape of the heart, but the utility of these prints is limited relative to VR visualizations because only the surface is physicalized. We foresee that as 3D printers become even more common and cost-effective, printing additional data on top of these surfaces will prove to be a helpful means for better comprehension of complex, multivariate, spatial data like those that our collaborators require. Finally, we note that using physical rendering processes to physicalize scalar spatial data combined with surface geometry data have the potential to make these data more graspable.

chapter four

Creating Spatial Data Visualizations with Handcrafted Physical Media

4.1 Introduction 35
4.2 Related Work 36
4.2.1 Artistic Involvement in Data Visualization 36
4.2.2 Visualization Interfaces for Non-Programmers 38
4.2.3 Visualization Systems 38
4.3 Design Requirements 40
4.3.1 Artistic Requirements 40
4.3.2 Technical Requirements 41
4.3.3 Printmaking Metaphor 42
4.4 Implementation 42
4.4.1 Terminology 42
4.4.2 Architecture Overview 44
4.4.3 The Visualization Design Interface 47
4.4.4 Visualization Schema 51
4.5 Results 52
4.5.1 Design Sessions Enabled by the Architecture 52
4.5.2 Design User Interface Insights 54
4.5.3 Additional Results 55
4.6 Discussion 56
4.6.1 Implications for Collaborative Scientific Visualization 56
4.6.2 Implications for Artistic Expression 57
4.7 Conclusion 58

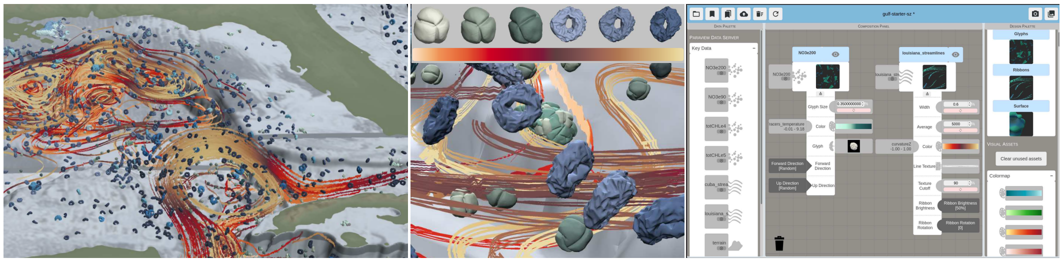

Figure 4.1: A visualization of biogeochemistry data in the Gulf of Mexico (left), a close-up view showing handcrafted elements of the visualization (center), and the visualization design interface used to make this visualization (right).

4.1 Introduction

[1]Scientists and society as a whole have a serious need for artist-designed visualizations which have a broad scope of visual vocabulary that enchants and enlightens viewers. The necessity for artistic involvement in the visualization community has been substantiated by numerous accounts, such as Cox's work in visualizing high-dimensional spaces [Cox, 1988] and Samsel's application of sophisticated artist-designed colormaps to highlight crucial parts of a dataset [Samsel et al., 2018a]. Figure 4.1 (*left* and *center*) show a small sample of the visual variety that is produced with an artist-designed visualization of biogeochemistry data in the Gulf of Mexico, produced in collaboration with climate scientists and computer scientists. However, achieving a high level of collaboration and communication between technical and artistic fields has been difficult in practice due to disparate vocabulary and differences in the process for investigating new information in each field. Most traditional visualization software focuses on the technical utilities of visualization, including data management pipelines and processing algorithms. While these utilities are quite valuable, they also have a significant influence on the design of the visualization software architecture they are encompassed in. Based on first-hand comments from artists involved in data visualization, the visual design tools within these data-first architectures are often limited in the visual vocabulary they can express [Samsel et al., 2019a], and their user interfaces can introduce an additional cognitive load which interferes with the creative process. In this chapter, we propose an alternative architecture which better supports artistic methodologies for creating visualizations by

[1] The latest version of the architecture and user interface is open source and available online at `https://github.com/ivlab/ABREngine-UnityPackage`. The visualization design interface in this chapter originated from work published at the 2020 IEEE VIS Arts Program entitled *Printmaking, Puzzles, and Studio Closets: Using Artistic Metaphors to Reimagine the User interface for Designing Immersive Visualizations* [Herman et al., 2020], and exhibited during a virtual workshop at the 2020 IEEE VIS conference [Keefe et al., 2020]. This work was accomplished in collaboration with members of the Sculpting Visualizations Collective, including Seth Johnson, Annie Bares, Francesca Samsel, Greg Abram, and Daniel F. Keefe.

using physical rendering processes that start in the real world.

In Chapter 3, we saw a physical rendering process that takes digital 3D multivariate spatial data and renders it using physical fabrication processes. The work in this chapter takes an inverse approach: we start from handcrafted, physical media elements created by artists and use them to produce visualizations with a broad scope of visual vocabulary that enhances the comprehension of data and enchants viewers. Artistic engagement in visualization shares many commonalities with data physicalization, including the use of physical materials and processes during the creation of a work. In physicalization, we generally interpret this need for physicality as a requirement to support the transferral of visualizations and visualization elements from the digital into the physical.

However, many artists prefer to begin the creative process in the physical world, working with physical materials and tools to produce their craft. This preference leads to new challenges and opportunities for computer science and data visualization research: How do we, as computer science researchers, meet artists halfway by enabling the use of physical materials and media to render data visualizations? How do we support non-linear, back-and-forth artistic processes for rapidly iterating through many visual designs for individual variables on a complex multivariate 3D spatial data visualization, and how do we enable artists to view their work with extended reality to see their 3D data in 3D? And ultimately, how do we support the "inverse" physical rendering process of encoding 3D multivariate spatial data with elements from the physical world? These are all still open questions in computer science and visualization research, and this chapter seeks a possible solution for supporting artistic processes and methodologies in data visualization by enabling the use of handcrafted physical media as visualization elements.

4.2 Related Work

4.2.1 Artistic Involvement in Data Visualization

Visualization has been a defining part of humanity's sense-making practices, ranging from ancient Mesopotamian tokens around 5500 B.C.E. [Schmandt-Besserat, 2013], to Leonardo da Vinci's work in visualizations that fused art and science, to the present-day IEEE VIS conference where prolific visualization researchers come together to advance the possibilities of the field. Artists have been involved in data visualization for as long as the field has existed, and have lent their creativity towards better understanding of and accessibility to complex data. One common thread throughout artists' engagement with data has been a preference for physical materials and processes, and when applied to data this is known as *data physicalization*. Physicalizations are a diverse medium, ranging from palm-sized sculptures showing world data as 3D printed crystals [Grossman, 2002], to park-sized installations at all-night art festivals centered on climate change [Swackhamer et al., 2017], to data-driven ice sculptures representing glaciers that make viewers gasp as they melt [Segal, 2015]. They have the potential to make data more comprehensible by leveraging both haptic and visual senses, and they have even been shown to be faster and more accurate than some digital visualizations [Jansen et al., 2013], and can create significantly more emotional engagement than digital visualizations using metaphor, physical material properties, and embodied interactions with the data [Wang et al., 2019].

Even when we cannot physically *touch* data physicalizations (i.e., if they are in a display case, or experienced through a digital intermediary), their benefits are still present through *imagined touch* [Peck et al., 2013]. Though imagined touch has been primarily explored in the online retail context, it has been increasingly applied to data physicalizations, with a recent physicalization survey noting that though the authors were not able to physically touch most of the works surveyed, they found themselves discussing the haptic qualities of the physicalizations anyway [Hornecker et al., 2023]. Although the multisensory *perceptual* benefits of physicalizations cannot be replicated through digital displays, with imagined touch we see that at least some of physicalization's emotional and embodied advantages remain. The work in this chapter build on this concept by using inverse physical rendering processes to encode data in digital visualizations with physical, real-world objects.

In tandem with artists' engagement in physicalization and data visualization, there have been technical developments to support creativity in the vast digital world. One research area in Human-Computer Interaction that has emerged from the drive to support creative methodologies with technology is Creativity Support Tools, which strives to remove as many barriers as possible for artists engaging with technology [Shneiderman, 2000], in essence "designing with low thresholds, high ceilings, and wide walls" [Shneiderman, 2007]. Common barriers that exist in technology for creatives are rigid and inflexible user interfaces, and lack of access to prior examples of work [Shneiderman, 2002]. Several solutions to these issues have been addressed, including use of metaphor and "fun" interfaces [Shneiderman, 2004], visual exemplars [Terry and Mynatt, 2002], and sketch-based interfaces, which are an entire research area of their own. Creativity Support Tools in visualization aim to not only to provide user interfaces that work for the creative process but to actively promote the discovery of new visual combinations that work to show data in a new light.

4.2.2 Visualization Interfaces for Non-Programmers

Sketch-based interfaces have emerged as a key entry point for artists and other creative individuals to get involved in visualization design. Artists have previously used sketch-based interfaces to prototype 3D visualizations [Keefe et al., 2005, Keefe et al., 2008a], create custom illustrations of 2D fluid flows [Schroeder et al., 2010], sketch free-form glyphs [Schroeder et al., 2010, Isenberg et al., 2008], and create multi-layered animated 2D visualizations of a variety of datasets [Schroeder and Keefe, 2016]. The visualization design user interface and architecture described in this chapter builds on prior work in "sketching user experiences" by supporting flexibility and affording rapid visual exploration of visualization designs [Buxton, 2010].

One no longer needs direct access to a supercomputer and need not be an expert programmer to engage with data [Viégas and Wattenberg, 2007]; many tools now exist that harness the power of traditional programming languages for visual design while providing built-in commands that abstract common graphics code, making it easier for non-programmers to write "creative code" to achieve powerful visual effects. One such tool is Processing [Greenberg et al., 2013], which is based on the Java programming language. Many art installations have used Processing for data visualization and background graphics, including *Orbacles*, where a Processing-based flock simulation was projected on the ground [Keefe et al., 2022a] and *city flows*, which uses Processing to visualize bike sharing data from three cities [Nagel and Pietsch, 2015].

Another approach to making computational visuals has been to avoid conventional programming languages and opt for an entirely visual experience for users instead. The Unreal game engine uses a node-based programming approach ("blueprints") which are in this style [Valcasara, 2015], and Blender's compositor and material editor use nodes to control the final look of a rendered piece of digital art [Vepakomma, 2014]. Another style of visual programming with puzzle pieces has been frequently used as a learning tool for computational problem-solving, for example with Blockly [Pasternak et al., 2017] and Scratch [Resnick et al., 2009]. The visualization design interface in our architecture melds these two approaches, employing a puzzle-piece approach to building visualizations.

4.2.3 Visualization Systems

State-of-the-art visualization software systems have powerful data management utilities for scientists and technologists, but their user interfaces fall short of supporting artistic workflows well. This is not a coincidence – writing a piece of software that is up to "artist spec" is notoriously difficult [Buxton, 1997], and for many visualization system developers, the artist spec is not a priority. We push back against this set of priorities because while the data processing utilities are a crucial part of the process, the end result is a *visual*ization, meaning the process for visual design should be given careful consideration in the structure of the software and its user interface. Giving this additional consideration becomes difficult because many modern visualization applications (i.e. ParaView [Ahrens et al., 2005] and VisIt [Childs et al., 2012]) are constructed as monolithic, self-contained applications. These applications do support nonlinear design processes to a certain degree, but to fully support artistic processes in software, modifications to their user interfaces are required to make them more flexible. The code for the user interfaces of these applications is technically modular and separable from the data management code, but in practice the UI toolkits these applications are built on are still too rigid to fully support creative methodologies. While there has been prior work emphasizing the "visual" in visualization software by building custom code for the visualizations and their design, these applications are limited in their generalizability between datasets. Some applications are built around a specific dataset (e.g. *city flows* [Nagel and Pietsch, 2015]), and others are restricted to 2D data (e.g. Drawing with the Flow [Schroeder et al., 2010] and Visualization-by-Sketching[Schroeder and Keefe, 2016]).

The Artifact-Based Rendering (ABR) technique [Johnson et al., 2019b] uses more generalizable approach than previous work by using an integrated data pipeline from ParaView [Ahrens et al., 2005], a powerful, parallel-computing wrapper for the Visualization Toolkit (VTK) [Schroeder et al., 2004]. ABR supports artists in the creation of complex visualizations with 3D, multivariate, spatiotemporal data by allowing the use of handcrafted artifacts as elements in a visualization. Stage 1 of the ABR technique starts in an artist's studio with clay, ink, paper, and other forms of traditional media. During Stage 2 of the ABR technique, artists' creations (sculpture, ink washes, chine collé, etc.) are digitized using 2D and 3D scanning techniques, resulting in visualization assets (*VisAssets*) that are usable by the ABR visualization engine. Lastly, these VisAssets are used to encode variables a dataset, leading to visualizations with a decidedly handcrafted aesthetic [Samsel et al., 2019a]. The ABR technique forms the basis for the architecture presented in this chapter, but it should be noted that there are marked differences between the "ABR Engine" mentioned in this chapter and the "ABR Technique" showcased in the 2019 IEEE VIS paper. The most notable difference is that the visualization system used in the original paper still used a monolithic architecture; both the graphics and user interface for creating visualizations were tightly coupled to the data processing pipeline. As such, it still fell short of the "artist spec"; in practice requiring a computer scientist to "Wizard-of-Oz" the visualization controls while the artist made the design decisions.

4.3 Design Requirements

In visualization system building research, it is common practice to start with the data specification and move towards a final visualization. This approach elegantly succeeds for data processing tasks like creating isosurfaces, sampling glyphs from volumetric data, and seeding streamlines; however, it also leads to visualization architectures that are less inclined towards artists' creative methodologies. We take a different path. We bring the visualization design goals of our artist stakeholders to the forefront of the software development process, all the way from early mockups to the current implementation. Working closely with artists throughout the entire development process of the architecture allows us as software developers to focus on making the software accessible to humans while still retaining its entire technical capacity. In this section, we identify a number of essential design requirements for a visualization software architecture outlined by our artistic collaborators, the corresponding technical requirements this architecture must support, and a means of encoding these technical requirements using metaphor.

4.3.1 Artistic Requirements

Rapid, nonlinear exploration of visual design space

Artists working in their studios employ a nonlinear process which involves repeatedly making detailed changes in a piece of artwork and taking a step back to observe the work as a whole. This type of process encourages a creative mindset where artists can lose themselves in the work and instinctively rely on their training in visual principles. When applied to data, the process results in unique visualizations that can help stakeholders understand their data better [Johnson et al., 2019b] and are clearly influenced by each artist's individual style [Samsel et al., 2019a].

Traditional-media artifacts

A big contribution to the individual style of artist-designed visualizations is the use of artifacts created with traditional, physical artistic media. Such artifacts include hand-sculpted clay glyphs and colormaps based on other artwork (Figure 4.1 *center top*), as well as hand-drawn lines and textures (Figure 4.10 B and C). Incorporation of these artifacts is essential to fostering the creative processes of artists working with data [Samsel et al., 2019a], and the physical artifacts afford unique opportunities to leverage *imagined touch* on the final visualization.

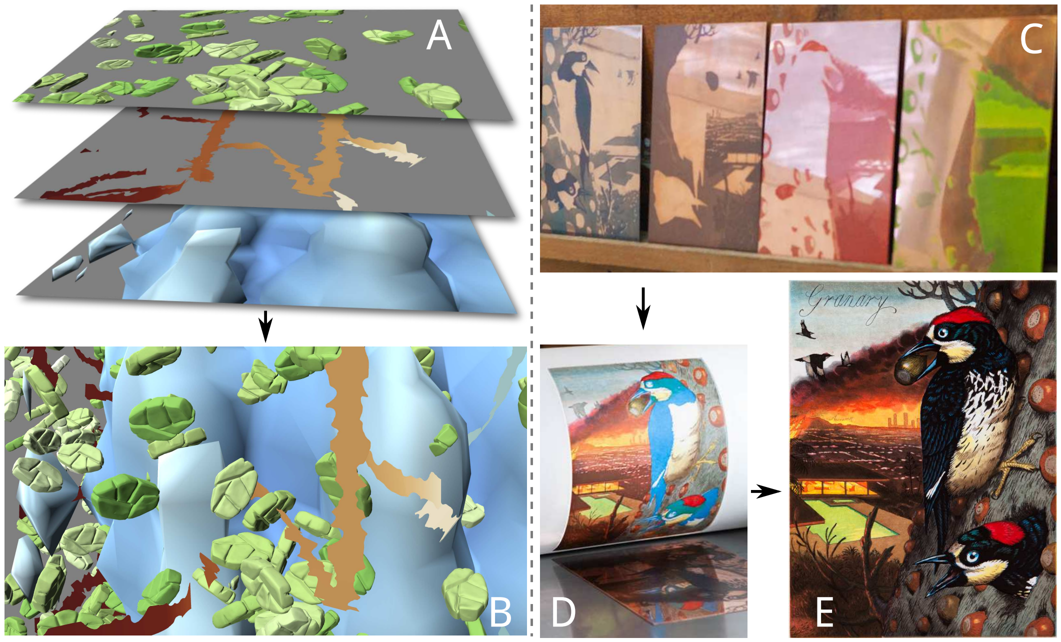

Figure 4.2: In computer graphics visualization, every visualization layer (A) is composited together into the final image (B). In intaglio printmaking, plates (C) are covered with colored ink and run through a printing press (D) to obtain impressions, many of which are combined into the final edition (E). *Right images Copyright 2020 Walton Ford and Wingate Studio; used with permission.*

4.3.2 Technical Requirements

Flexible user interfaces

Requirements 4.3.1 both necessitate a visualization design interface that is suited to creative processes. In the style of Creativity Support Tools [Shneiderman, 2000], this means supporting artists in their exploration of the visual design space by enabling them to make rapid changes to the visualization and receive immediate visual feedback.

Real-time visualization control

To achieve the prior requirements, the visualization software architecture must support:

A. Quickly swapping visual elements and data

B. Modifying visual elements based on data

C. Viewing the visualization from various angles

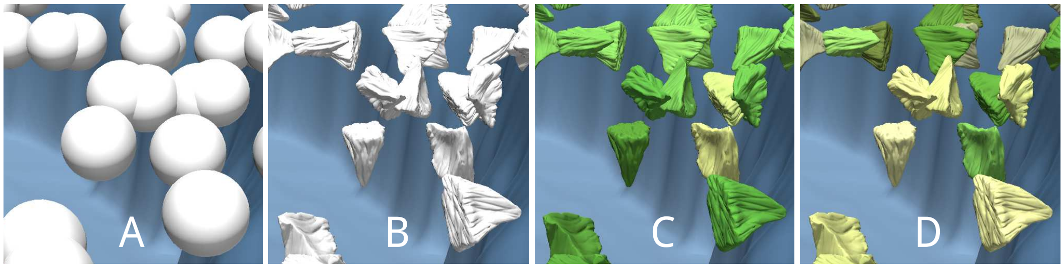

Figure 4.3: The construction of a glyph data impression, showing (A) applying the chlorophyll point sampling key data, (B) applying a glyph from the design palette, (C) applying a colormap and *Salinity* variable, and (D) applying the *Temperature* variable with the same colormap. Notice how the glyphs remain in the same location throughout because they all use the same Points key data object.

4.3.3 Printmaking Metaphor

Effective visualization software for artistic stakeholders finds a balance to achieve all the technical and artistic requirements, often through the use of metaphor; this is another essential requirement for the architecture. Intaglio printmaking was the metaphor upon which the architecture is built. Printmaking helps bridge the gap between the technical requirements of artist stakeholders and their technical implementations because it bears a resemblance to the way computer graphics are drawn and is a process that a majority of artists are familiar with. Figure 4.2 shows a step-by-step comparison of the printmaking process with the computer graphics rendering process.

4.4 Implementation

In this section, we present the design and implementation details of the new data visualization architecture that supports artist-focused visualization design interfaces. The entire software design process surrounding the architecture is focused on providing artists with a familiar means of creating a piece of visual art – in this case, a multivariate visualization. Thus, developing a visualization design user interface that fits well with artists' nonlinear creative processes was at the forefront of the team's minds as we created the architecture. To this end, we enumerated and deconstructed the elements of modern visualization software and put them back together into a new architecture that better supports artists' creative practice.

4.4.1 Terminology

The architecture borrows terminology from both intaglio printmaking and traditional data visualization. Terms from the ABR technique are also used, including VisAssets,

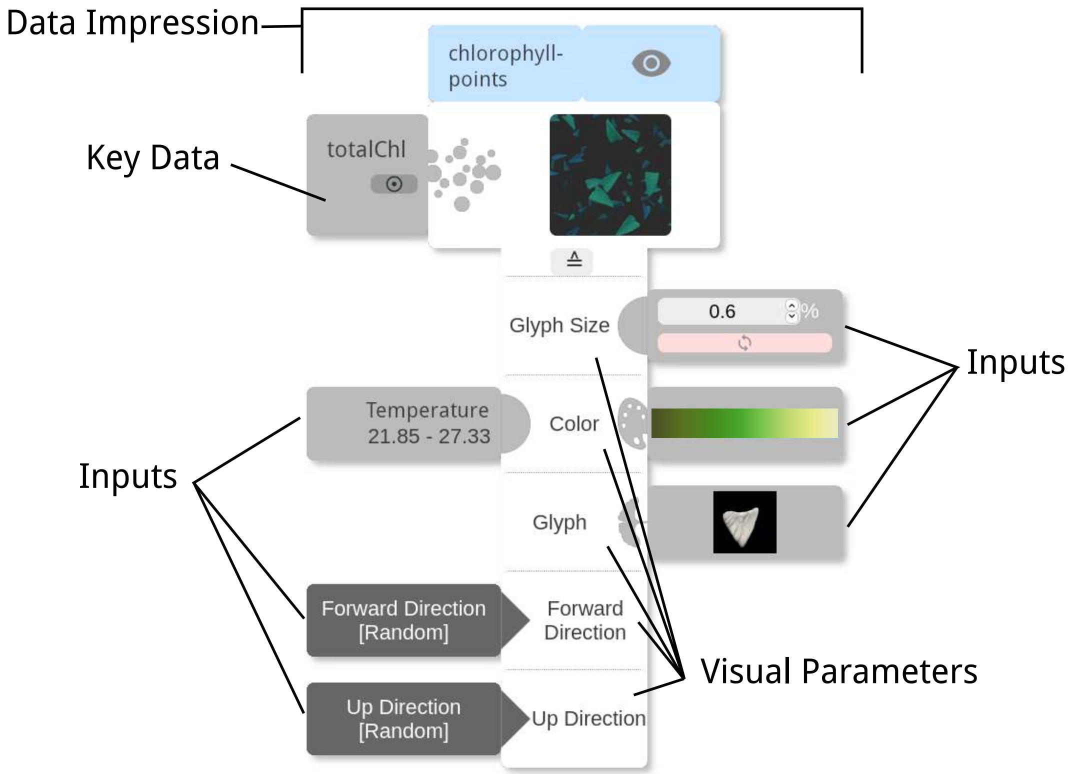

Figure 4.4: A Glyphs data impression (visualized in Figure 4.3D) with inputs populated by an artist.

the artist-created visual elements that can be used in a visualization (shown in Figure 4.1 *center top*). In printmaking, the plate with a registration (or the "structure") in printmaking is called a *key plate*, and each pass through the printing press is called an *impression*. These artistic terms are adapted for use in the data visualization world, hence we get *plates*, *data impressions* and *key data*. Key data are the geometric structures upon which we apply our designs to – *Points*, *Lines*, and *Surfaces*. One example of such a structure is the Points key data object with a glyph sampling of chlorophyll in the Gulf of Mexico shown in Figure 4.3, where glyphs are denser with higher chlorophyll concentrations and sparser with lower concentrations. Plates specify the available inputs on a particular type of data impression; this defines what type of key data they accept and other inputs like *Glyph Size* (numeric), *Color Variable* (scalar variable), and *Pattern* (VisAsset). Individual data impressions ("towers," shown in Figure 4.4) are initialized from plates, and are comprised of one or more key data objects along with all the visual designs applied to them. For instance, we might take the approach shown in Figure 4.3 and first apply a hue-varying green colormap, then apply a handcrafted triangular glyph to the chlorophyll point sampling.

The new metaphoric vocabulary is augmented with standard visualization vocabulary like *variable* and *dataset* to create a unified terminology in the architecture. We consider a dataset to be a collection of related key data objects, each containing variables. Variables may be present in more than one key data object. For instance, the *Temperature* scalar variable may be present in both the chlorophyll point key data and the ocean current streamlines key data, and the vector variable *FlowDirection* may be present in both of these key data as well.

4.4.2 Architecture Overview

In software systems, modularity is generally considered "good design", and has precedent in many systems today [Schwanke, 1991]. To achieve modularity and maximize the separation of the visual design tools from the data management and rendering utilities, the team embraced a distributed architecture where each component can run on its own computer, and they communicate across a network. Decoupling the visual design tools from the data management utilities was the first step towards achieving a modular, distributed architecture. The architecture is separated into four main components:

- A web-based design user interface
- A web server
- A rendering engine
- A data host

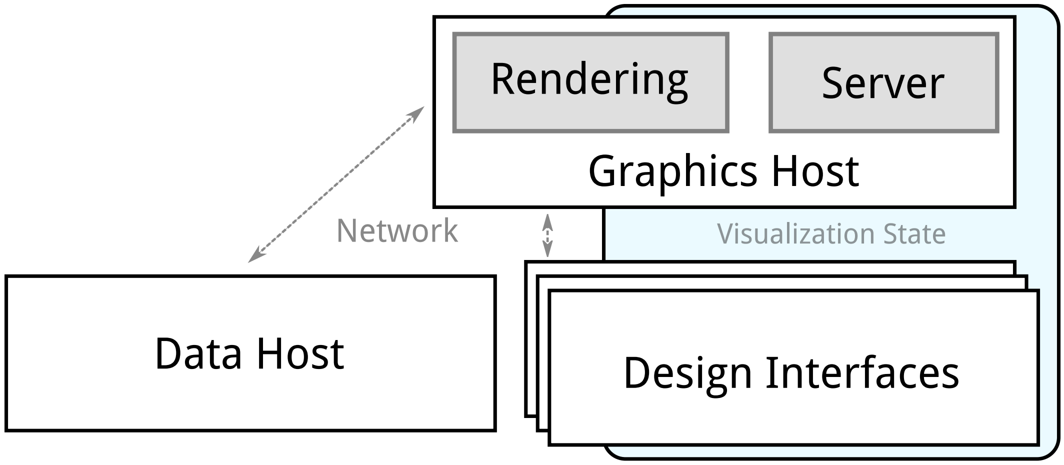

Figure 4.5: Data sources hosted on a supercomputer, data visualization graphics and a web server hosted on a powerful graphics computer, and visualization design interface clients all work together to allow artists to build visualizations of 3D multivariate data which are rendered at real-time frame rates. Visualization state is shared between the rendering engine, the server, and the design user interfaces.

Figure 4.5 depicts the components of this architecture and their relationships over a network. An artist creates the visualization in a web browser using the design UI (Figure 4.6), which is served from the web server on the graphics host computer. The server acts as an intermediary between the design UI clients and the rendering engine. A powerful graphics computer runs the rendering engine and displays the final visualization at interactive frame rates. The computer on which the data resides is known as the data host, which also supplies the data management utilities and sends preprocessed geometric data to the rendering engine. For the purposes of concrete demonstration, the examples shown in this chapter use a combination of the Artifact-Based Rendering technique [Johnson et al., 2019b], ParaView [Ahrens et al., 2005], and a Python Django server to implement the components of the architecture, but each of these components may be swapped out at any time.

Rendering engine

The ABR engine is the rendering engine in this implementation of the architecture, which uses the Unity game engine, C# scripts, and ShaderLab shaders to render the artist-driven visualization graphics. The engine supports creative methodologies in visualization by allowing artists to use their own work, often created in physical media, and is connected with a web-based visualization design user interface (Figure 4.6). The rendering engine accepts commands in real-time from the visualization design interface, supporting Requirement 4.3.2A. As artists construct a visualization using the interface, states can be saved to JSON files at the artist's discretion so they can come back to work on it later or share the state with others. If the rendering engine is running on a computer with a head-mounted Virtual Reality display connected, the resulting visualization can also be displayed in VR at interactive frame rates. The rendering engine also supports Requirement 4.3.2C by supplying standard 3D navigation controls in the visualization display window.

Data host

To obtain the data, the rendering engine connects through sockets to the data host. In the data host, volumetric data are preprocessed into geometric forms (points, lines, and surfaces), then further processed in the rendering engine to get the renderable meshes. Glyphs are passed through an instance-rendering shader, lines are rendered as a quad strip, and surfaces are rendered as a series of triangles.

Web server

On the user-facing side of the architecture is the web server which handles all connections and messaging with the design user interface. The server is run within the same file system as the rendering engine so that data and artist-created VisAssets can be shared between these components. For instance, the design UI fetches the colormap thumbnails directly from the server, and the rendering engine uses colormap files in the same directory to apply to the visualization. The server also passes messages between the design UI and the engine using WebSockets [Fette and Melnikov, 2011]. All messages between the design UI and the ABR engine are in JSON format.

Visualization design interface

The visualization design interface is built for web browsers in order to support as many devices as possible, and to take advantage of the breadth of modern web libraries available. Specifically, the interface uses jQuery UI for most of its primary interactions, including dragging and dropping puzzle pieces. The interface also includes a colormap editing utility (shown in Figure 4.8) that takes advantage of the D3 visualization library [Zhu, 2013] for showing the histogram of the variable attached to a particular colormap. Because it is implemented as a web app, the design interface can be run on a different computer from the visualization – even on a tablet across the country.

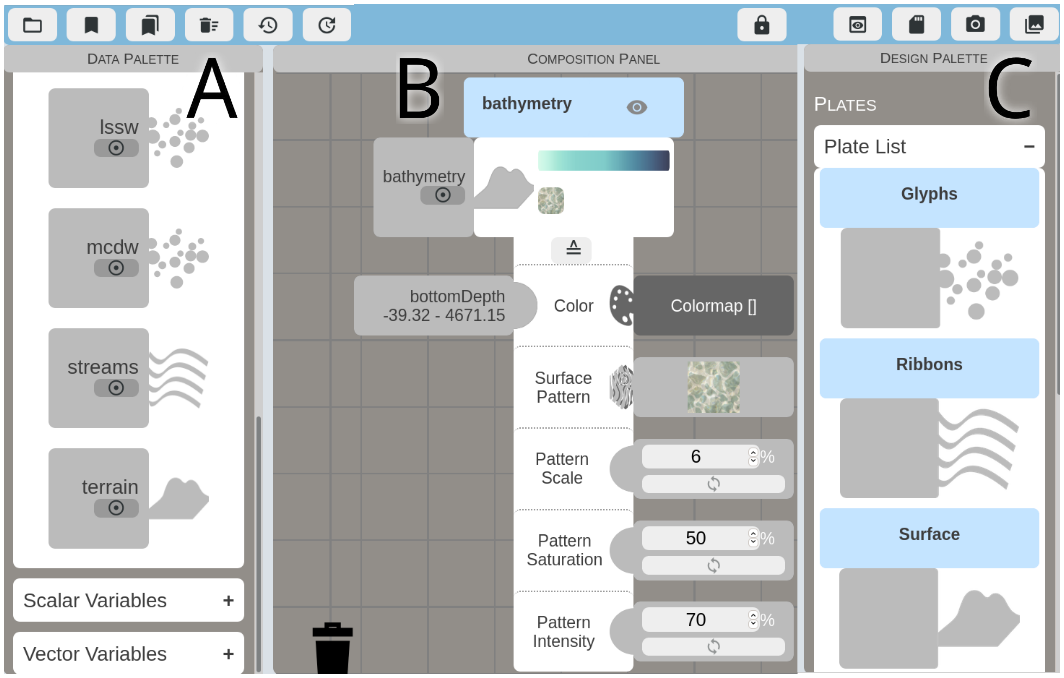

Figure 4.6: The design user interface, showing the data palette (A), the composition panel (B) with one data impression, and the design palette (C).

4.4.3 The Visualization Design Interface

Figure 4.6 shows the web-based 3-panel design interface that allows for rapid iteration on the visualization contents. When considering artists' need for fast, nonlinear iteration of visual designs (Requirement 4.3.1), the fundamental processes used in the construction of a visualization need to change from the traditional data-first approach. For instance, an artist may consider both the *Ribbons* and *Clay Forms* in Figure 4.7 to be "lines," and features like seamlessly changing from ribbons to clay forms is something most artists would expect from the software they work with. However, from a data-first computer graphics standpoint, rendering a ribbon (quad strip) is significantly different than rendering a series of instanced glyphs; in fact the easiest approach to make this kind of change is usually to make a new glyph layer, copy all the visual encodings from the old ribbon layer, and then discard it. The visualization design interface uses the metaphor of a "data impression" to convey the technical constraints of the computer graphics algorithms in a way that is consistent with artists' studio processes.

Data Impressions

Data impressions are the means with which artists create visualizations in the design interface. Each data impression is a single "layer" of the visualization, and artists build their visualization composition by creating many data impressions with different key data. Data impressions are initialized by dragging and dropping a plate from the design palette panel on the right into the composition panel in the center of the design user interface (see Figure 4.6). Figure 4.4 depicts a single Glyphs data impression which uses the chlorophyll point sampling from Figure 4.3 (D) and two artist-chosen VisAssets – a colormap and a glyph.

Data impressions have one or more key data. In Figure 4.4, the Glyphs data impression has a Points key data object attached to it. Like every variable and VisAsset in the interface, the key data puzzle pieces have evocative iconography showing an artist's interpretation of that type of key data. Key data are selected from the left data palette panel and dragged into matching data impressions in the center composition panel of the design interface (see Figure 4.6).

Additionally, data impressions each have a set of visual parameters that the artist can change. Each parameter may have one or more inputs, for instance in Figure 4.4 the *Color* parameter has two inputs: *Colormap* for the artist-created colormap VisAsset, and *Color Variable* for the scalar data variable to be encoded with that colormap. Similar to key data, scalar and vector variable puzzle pieces are dragged and dropped from the left data palette, and VisAssets are dragged and dropped from the right visual design palette into matching puzzle slots on a data impression in the center composition panel. Two methods are used to indicate puzzle piece compatibility – the iconography of the puzzle piece shape must match its input slot, and a subtle blue glow appears over each compatible slot when dragging a puzzle piece.

Building Visualizations

Using the interface, artists create visualizations in a manner similar to the processes they use in the studio. Data impressions provide a logical means of organizing components of a visualization, and inputs to these data impressions allow for quick, iterative changes to be made to a visualization at the drop of a puzzle piece (which addresses Requirement 4.3.2). We provide two concrete examples of working with the interface which showcase possible ways an artist may desire to create a visual outcome. Both examples use data from a Gulf of Mexico biogeochemistry dataset [Petersen et al., 2015], and have Surface data impressions of the ocean floor and land for context. The generic process for getting started with a first data impression in the design user interface is as follows:

1. Initialize a data impression from a plate
2. Assign a key data object

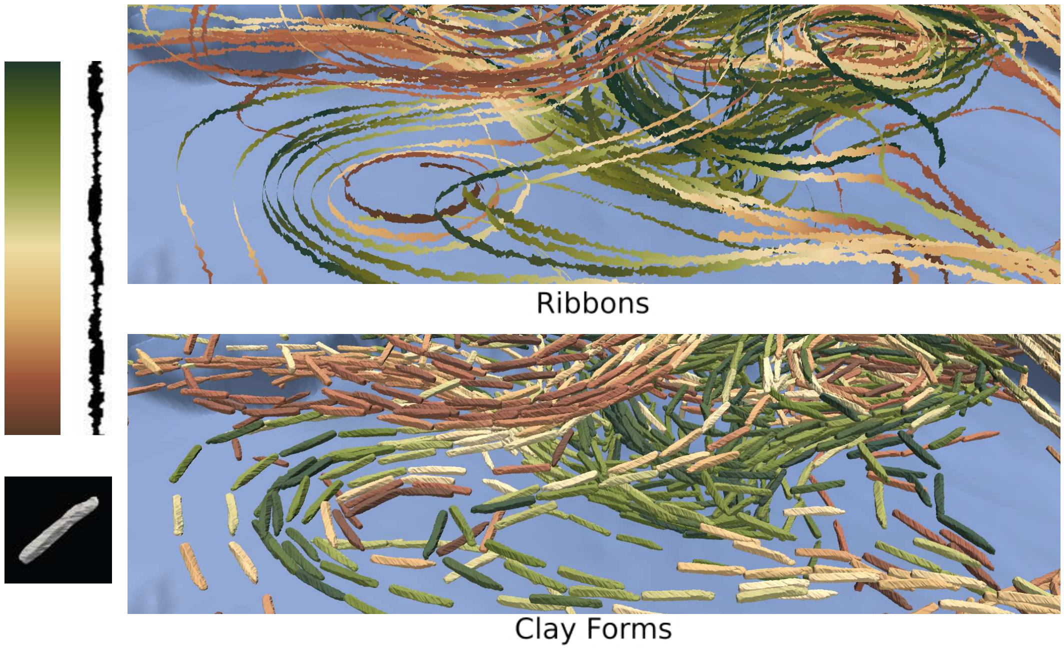

Figure 4.7: Two line impressions, one made with a ribbons plate and another made with a glyphs plate.

3. Assign VisAsset and variable inputs
4. Iteratively adjust visual parameters

First, let us return to the previously mentioned situation where an artist wants to change the type of line for the ocean current streamlines south of Louisiana shown in Figure 4.7. Let's say that an artist wants to begin with representing these currents with Ribbons. A Ribbons data impression is instantiated, the proper key data are applied, and the artist chooses VisAssets for the *Line Texture* and *Color* visual parameters. At this point, the artist may decide that they want to use clay forms instead of ribbons to visualize the ocean currents. Recall that the computer graphics algorithm for rendering clay forms is entirely different from that of the ribbons. This constraint is directly encoded within the design interface: the artist cannot drag the Points key data of the line onto the Ribbons data impression – they must create a new Glyphs data impression and copy across each data variable and visual element to the new impression.

The second example (Figure 4.3) shows a point sampling of chlorophyll density near the western coast of Florida. A Glyphs data impression is instantiated, and the chlorophyll point sampling key data object is applied to the impression. Next, a glyph VisAsset is chosen from the design palette and applied to the *Glyph* input, a colormap from the design palette is then applied to the *Colormap* input, and a scalar variable *Salinity* is applied to the *Color Variable* input. Lastly, if the artist decides to visualize

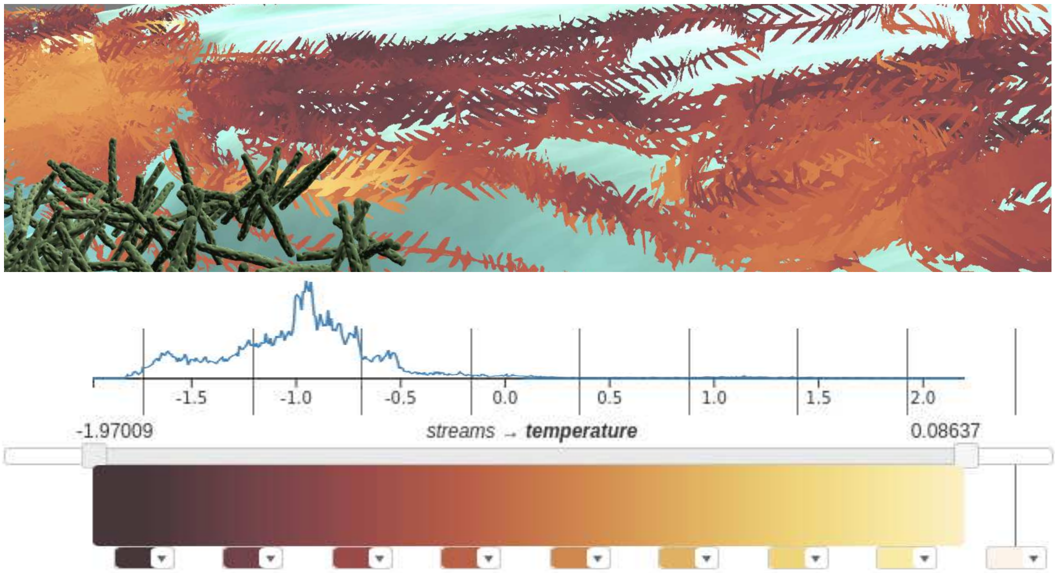

Figure 4.8: The colormap editor integrated with the visualization design user interface allows for on-the-fly changes to the colormap based on a histogram of real data values.

Temperature instead of *Salinity*, they need only drag in the *Temperature* scalar variable from the data palette to replace the *Color Variable* input.

Working with VisAssets

Artists to want to use their own work when creating visualizations, so any artist-focused interface for visualization design must keep this in mind (Requirement 4.3.1). The Artifact-Based Rendering technique allows artists to incorporate their own textures, colormaps, lines, and glyphs which affords the resulting visualizations a much greater visual vocabulary [Samsel et al., 2019a]. The design interface allows for two complementary strategies for working with VisAssets: downloading them from a curated library [Samsel et al.,], and working with local copies.

Handcrafted artifacts (glyphs, colormaps, lines, and textures) are curated by artists in the Sculpting Vis Collective, and a series of ABR web applets [Johnson et al., 2019b] are used to convert them into usable VisAssets. In order to take full advantage of ABR, artists will often need to modify these VisAssets while designing a visualization. The design user interface features a colormap editor (Figure 4.8) inspired by ColorMoves [Samsel et al., 2018b] which allows the artist to create custom, dataset-specific colormaps and view the resulting color changes on the data in real-time.

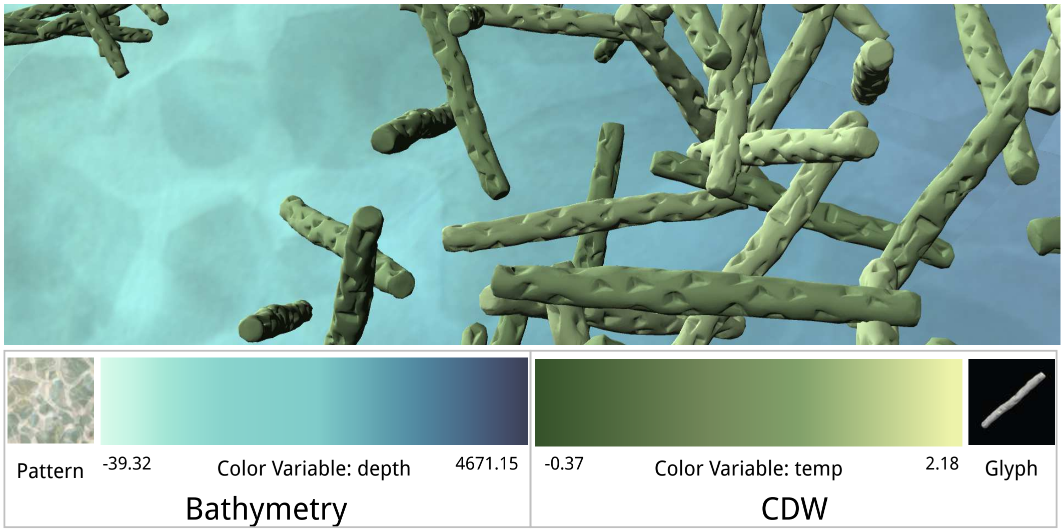

Figure 4.9: ABR visualization created by artist Francesca Samsel with a Glyphs data impression showing Circumpolar Deep Water (CDW) and a Surfaces data impression showing the Bathymetry below the Ronne-Filchner ice shelf, described by the JSON state in Listing 4.1 at the end of the chapter.

4.4.4 Visualization Schema

We add an additional component to augment the architecture presented so far – a visualization state specification based on a JSON Schema [Pezoa et al., 2016]. Schemas have a long-standing precedent for use in defining a visualization state; Snap-Together Visualization [North and Shneiderman, 2000] uses a relational database schema and Vega-Lite [Satyanarayan et al., 2016] uses a JSON schema to define interactive visualizations. The schema describes a formalization of an artist-created visualization state, such as that shown in Figure 4.9 and Listing 4.1 at the end of the chapter. The following describes the roles of this schema and their impact on the flexibility and stability of the architecture.

Schema as a description of capabilities

The schema defines the plates that are available in the architecture. This allows developers to make algorithmic improvements to the rendering engine without affecting code in the user interfaces or destroying backwards compatibility with visualization states. The schema also provides definitions to every input so that the rendering algorithms can be improved or swapped out with complete transparency to the artist.

Schema as a contract

The schema also ensures that the visualization state representation is consistent across the architecture. The visualization state schema has four main components: data impressions, local VisAssets, scene information, and UI data. The most important part of the visualization state, data impressions, is a list of artist-created layers. Each data impression has a unique identifier, and its inputs are populated with the VisAsset identifiers, data paths, or primitives that the user has dragged and dropped in the artist-facing design interface. Other scene data such as lighting information is also initialized and adjusted through the interface. In addition, the state also stores VisAssets the artist has modified, as well as any UI-specific data that may need to be saved for a particular interface, such as the location of the data impression towers in the composition panel.

4.5 Results

The architecture enables new remote multivariate 3D visualization collaboration in a way that was not previously feasible. Another important result for the architecture was the feedback and insights on the interface and visualization design provided by artists who used the interface. In this section, we describe the remote design sessions that were conducted with a variety of geographically distributed collaborators, as well as several insights from the artist stakeholders who used the interface.

4.5.1 Design Sessions Enabled by the Architecture

The network-based implementation of the architecture allowed our team to conduct design sessions with people across broad distances. During each design session, the participants used the web-based visualization design interface to build visualizations of curated datasets such as biogeochemistry in the Gulf of Mexico [Wolfram et al., 2015] and ocean currents under the Ronne-Filchner ice shelf of the Antarctic [Petersen et al., 2019]. As necessitated by the COVID-19 pandemic, all design sessions were run entirely remotely via the networking and web capabilities of the architecture in conjunction with video conferencing software.

In the architecture's infancy, we facilitated two remote design sessions with artists. Both artists watched a short training video on the use of the interface prior to their sessions, and could ask as many questions as they wanted while creating a visualization. The first session was conducted with artist Deborra Stewart-Pettengill, who works with many forms of traditional artistic media and had no prior experience working with 3D scientific data. Stewart-Pettengill has lately been working with chine collé, a printmaking technique that involves thin pieces of colorful paper which are cut out and applied to

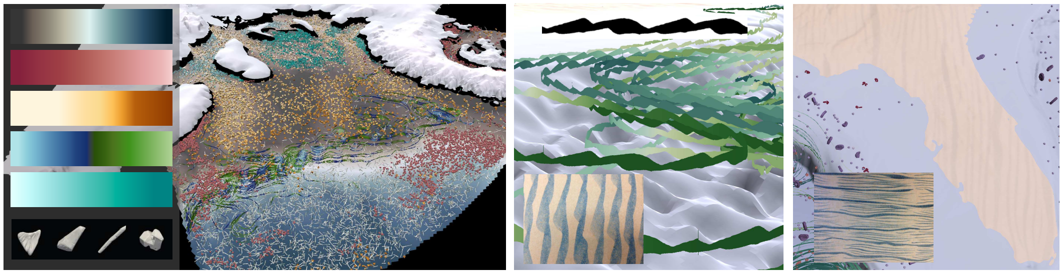

Figure 4.10: Francesca Samsel uses custom colormaps and her own hand-sculpted clay glyphs in this visualization of ocean currents mixing under the Ronne-Filchner ice shelf in Antarctica (A). Deborra Stewart-Pettengill commonly works with patterns, which have been digitized to form streamlines (B) in the Gulf of Mexico visualization and a texture on the land (C).

the print when it is run through the press. The interaction with the visualization design interface allowed Stewart-Pettengill to use her chine collé work in a visualization of biogeochemistry data in the Gulf of Mexico (see Figure 4.10 B and C). During the session, Stewart-Pettengill iterated through a total of 12 glyphs, 7 lines, 8 textures, and 26 colormaps from the Sculpting Vis library. Three of the lines were processed into VisAssets from Stewart-Pettengill's chine-collé work, which took an ABR expert approximately thirty minutes total to accomplish before the session started. Additionally, all 8 textures were sourced directly from the chine-collé work, which took an ABR expert about 10 minutes total to convert into VisAssets. Session one took approximately an hour and a half total.

A second design session was held with artist Stephanie Zeller, who has worked with 3D scientific visualizations previously including custom colormap creation. Through the use of the Sculpting Vis web applets [Samsel et al.,] and the colormap editor in the design user interface, Zeller created several new colormaps as well as making dataset-specific modifications to colormaps from the Sculpting Vis library to benefit the clarity of the variable they are applied to. During the session, Zeller used 14 glyphs, 9 lines, and 9 textures, and 56 colormaps. She edited 18 of these colormaps during the session using the colormap editor, and she created two colormaps from scratch using the Sculpting Vis web applets [Samsel et al.,]. See Figure 4.1 (*left* and *center*) for a view of Zeller's visualization, featuring her custom colormaps in combination with other VisAssets from the Sculpting Vis library. Session two was conducted in two parts, the first about one hour long and the second nearly two hours.

A few weeks later after several improvements were made to the architecture and user interface based on feedback from the earlier artist sessions, another design session was conducted with 7 members of the Sculpting Vis Collective. This session was focused on the remote collaboration possibilities enabled by Johnson's approach to remote graphics rendering [Johnson, 2020]. Each member of the design session viewed the visualization remotely through a VR headset or a desktop viewer if they didn't have access to a head-mounted display. During the session, the team was able to produce insights about the 3D multivariate data we were working with while one of the artists on our team worked rapidly on the visualization using the design interface.

At the IEEE VIS conference in October 2020, the Sculpting Vis Collective hosted a virtual tutorial on Artifact-Based Rendering which used the architecture to facilitate collaborative remote visualization design [Keefe et al., 2020]. During the tutorial, participants had the opportunity to learn about the philosophy of our team, how to use the software, and design recommendations for creating hand-crafted visualizations. Participants were then encouraged to practice with the visualization design user interface with the assistance of Sculpting Vis team members. Several example visualizations were given with the data the participants were to use, including the Antarctic visualization shown in Figure 4.10A. After this, participants were grouped with members of our team via video conference breakout rooms to expedite answers to questions. Nodes on a supercomputer were used to run 8 instances of the architecture for the tutorial. Each instance contained a data host, a rendering engine, and a web server for the design user interface. Tutorial participants connected to the design UI through web browsers on their personal computers and designed visualizations together with each other and the Sculpting Vis team via video conferencing software. The tutorial had approximately 40 participants total, about 10 of whom remained for the interactive portion to use the design interface.

4.5.2 Design User Interface Insights

From the beginning of the software design and development process, we worked with artists in our interdisciplinary Sculpting Vis Collective to determine requirements for the architecture and interface as well as enhance their functionality and form. During the formulation of the interface, the Sculpting Vis Collective held semiweekly design meetings that focused around next steps for developing the interface's capabilities as well as higher level discussions about the implications the interface and architecture will have on the visualization field. At each meeting, artists on the team discussed what was working well with the interface and what needed to change to better serve their creative processes. Even early in its evolution, the artists were excited about the possibilities that the interface allows for them while creating visualizations, including the rapid iteration through the visual possibility space and the scope and fidelity of visual vocabulary granted by the artists' use of their own media in the visualizations

(Requirements 4.3.1).

The team received further valuable insights during the artist sessions which had a significant influence on the look, feel, and capabilities of the user interface and architecture. One noteworthy item during these sessions was the artists' enthusiasm for using their own work in the visualizations; this reinforced Requirement 4.3.1 as an essential item for the architecture to support. When Stewart-Pettengill first saw her 8 chine-collé textures in her palette of the design user interface, she immediately recognized them: "Oh yeah, there are my strokes! Cool!" She went on to apply several of these textures to the lines depicting ocean currents as well as the surface depicting the landmass surrounding the Gulf of Mexico (see Figure 4.10 B and C). Zeller also experienced a similar excitement for using her own colormaps, applying these custom colormaps to nearly every variable in her visualization. Another piece of insight was observing how the artists worked with the interface; the importance of Requirement 4.3.1, supporting an iterative design process, was affirmed by the artists. Both artists approached the visualization design using a nonlinear process – each design decision came as a reply to the current state of the visualization. Zeller found this process similar to that which she uses in her large-format paintings; she works on minute details up close and repeatedly steps back to look at the "big picture" to make sure everything is visually balanced. Stewart-Pettengill also remarked that she is "just so used to experimenting with things" as she works, and that the interface enabled a similar experimentation process during her session.

4.5.3 Additional Results

Beyond the success of the above design sessions, we have observed other noteworthy results from our experience with the artist-focused design user interface and architecture.

Firstly, the software holds up to usage not only as a design tool but a means to interactively visualize data the way scientific sensemaking processes necessitate. The software has been used in two iterations of the Visualization course at the University of Minnesota. Students begin the course by handcrafting visual encodings from physical materials; the first project involves designing 20+ gradients that could convey different types of data (linear, divergent, discrete, continuous). Then, students used the design interface and architecture to create visualizations with real, multivariate 3D spatial data – for example, a time-varying dataset on the wildfire phenomenon "vorticity-driven lateral spread", from the 2022 SciVis contest at IEEE VIS[2]. After successfully designing a multivariate visualization using visual encodings they handcrafted, the students created several different interactive, queryable visualizations using their visual encodings as well,

[2]SciVis 2022 Contest: https://www.lanl.gov/projects/sciviscontest2022/

and many opted to use our software for their final project. While the final visualizations were not as diverse and meaningful as we had initially hoped, but we attribute this more to the fact that the majority of the classes were undergraduate computer scientists who lack significant experience in utilizing color, line, texture, and form to design effective visualizations, rather than software limitations; the students who truly engaged with the software and principles behind the inverse physical rendering process were able to create truly unique visualizations.

So far in the lifetime of the Artifact-Based Rendering technique and this architecture, we've also observed some unique behaviors from visualization viewers that indicate there's something special to the approach of creating visualizations with handcrafted elements from the physical world. From the original ABR paper [Johnson et al., 2019b], we observed a climate scientist looking at his biogeochemistry data and referring to the glyphs representing chlorophyll and nitrates:

> *They look like biology more than they look like plastic. They could be real, produced by nature. I think that people are going to underestimate that. At first, I'm perturbed, these don't look like plastic, then I realize this is not a problem but a major benefit.*

During many demos in our lab, people have reached out with their hands to try and touch parts of a VR visualization created by one of our collaborating artists with physical elements from her studio and this software, and have commented on how certain parts of the visualization look like "Cheerios" or like "popcorn". We hypothesize that individuals viewing a visualization created with our inverse physical rendering process are able to leverage imagined touch on the visualization because its elements originate in the physical world. That is, people are able to better relate to these handcrafted, physical-looking elements than geometrically defined, sterile primitives provided by most other current visualization software because they are initially created from physical materials, and engage the same "need for touch" that physical objects provide [Peck and Childers, 2003].

4.6 Discussion

4.6.1 Implications for Collaborative Scientific Visualization

Our distributed architecture provides the modular separation necessary to promote the use of artistic processes while working with big data, and allows for the flexibility to host different components on geographically distributed computers. This flexibility is evidenced both by the diversity of use cases that our team has already explored using this architecture.

In total, we held four major visualization design sessions with the architecture since its inception, all of which took place during the COVID-19 pandemic. While remote collaboration wasn't an initial requirement of the architecture's software design, it quickly became necessary to the ongoing progress of the project. Losev et al. describe a remote collaboration strategy "simulating a co-located design space" in which team members each have two webcams - one which to show the member's face and one to use showing sketches [Losev et al., 2020]. Our architecture employs a similar strategy, except instead of using a second webcam, the "sketching" and ideation on a visualization is done directly in the web-based design user interface. This is the strategy which was used during all of our design sessions thus far, and it demonstrated its effectiveness for generating unique and rich visualizations like those shown in Figure 4.1 and Figure 4.10. We anticipate that in the future, this type of remote collaboration will be incredibly useful for geographically distributed teams like the Sculpting Vis Collective.

Achieving accessibility to visualization design software at varying levels of graphics hardware is another benefit to the modular architecture. Since many actively-studied datasets are very large, it is not feasible to move them between computers; this has already been addressed with systems like ParaView [Ahrens et al., 2005] which allow the user to connect to another instance of ParaView running on a different computer. Our architecture provides this, in addition to two further degrees of accessibility – remote design and remote rendering of visualizations. Visualizations can be designed remotely via the artist-focused design interface, which only requires a device with a web browser and a network connection. Remote rendering allows the user to view their visualization designs without having a significant rendering capacity on their devices [Johnson, 2020]. These two techniques are steps towards lowering the barriers for artists working with data; the hardware requirements for designing a visualization are significantly reduced when using this architecture as long as there is a data host and rendering engine available via a network.

4.6.2 Implications for Artistic Expression

The speed at which one can explore visual alternatives with the design interface stimulates the artistic imagination of those who use it. All the artists who have used the interface so far have remarked that it was a fun experience for them. They also mentioned that working with the interface was very similar to their studio processes; it allows for rapid experimentation of visual alternatives as well as "taking a step back" to see the bigger picture of the visualization, which are all core principles in creative methodologies.

The interface's proclivity towards visual exploration was also evidenced by the number of VisAssets used by each artist during their design sessions. The interface removes the barriers to achieving this creative state so that artists can use their skills built up over a lifetime of experience to add their voices and visions to the multidisciplinary efforts to wrangle increasingly large and complex scientific data.

The visual variation exhibited by artists using the design interface in the new architecture was significant because its workflow enables each individual to contribute their artistic vision to the visualization design. Past work with the Sculpting Vis Collective (e.g. [Samsel et al., 2019a, Johnson et al., 2019b, Samsel et al., 2018a]) and experiences with the design interface so far indicate that a crucial part of artists interacting with data is using their own handcrafted elements in their visualizations. The architecture supports this aspect of artistic engagement in visualization, as evidenced by artists' use of their own chine-collé work, custom colormaps, and hand-sculpted clay glyphs.

4.7 Conclusion

This chapter introduced an "inverse" physical rendering process through which is now possible for artists to handcraft elements in the physical world and link them with complex, multivariate 3D spatial data, and finally view the visualizations in extended reality. The inverse physical rendering process has the potential to enable scientists to encode more variables more effectively in their visualizations to better comprehend their data. Throughout the creation of this architecture to support artists building 3D multivariate data visualizations, we have decoupled the components of current visualization software, developed user interfaces to support creative processes, connected these interfaces together with visualization rendering engines, and built a formalization of an artist-created visualization state as a JSON schema. The architecture introduces a distributed approach to visualization design software, ensuring that the specification for the rendering system is separate from the design requirements of any artist-facing user interface. Similar to previous work in artist-focused systems for visualization design, the emphasis of our work is on flexibility, both in a user interface that supports rapid iteration and in a modular architecture that allows easy modification of its software components. The Sculpting Vis Collective continues to work towards our common goal of providing artists in visualization with the best possible experience for the benefit of science and humanity at large, and the distributed architecture this chapter presents is a significant step forward for artists designing 3D multivariate visualizations.

```
{ "impressions": [ {
      "uuid": "12598986-819c",
      "plateType": "Surfaces",
      "name": "Bathymetry",
      "inputValues": {
        "Key Data": {
          "inputType": "PointsKeyData",
          "parameterName": "Key Data",
          "inputGenre": "KeyData",
          "inputValue": "/Ronne/key-data/bathymetry" },
        "Color Variable": {
          "inputType": "ScalarVariable",
          "parameterName": "Color",
          "inputGenre": "Variable",
          "inputValue": "/Ronne/scalar-vars/depth" },
        "Colormap": {
          "inputType": "ColormapVisAsset",
          "parameterName": "Color",
          "inputGenre": "VisAsset",
          "inputValue": "f31cfdb3-0f2f" },
        "Pattern": {
          "inputType": "PatternVisAsset",
          "parameterName": "Pattern",
          "inputGenre": "VisAsset",
          "inputValue": "285a2582-48f6" } }
    }, {
      "uuid": "99e2b165-2e12,
      "plateType": "Glyphs",
      "label": "CDW",
      "inputValues": {
        "Key Data": {
          "inputType": "PointsKeyData",
          "parameterName": "Key Data",
          "inputGenre": "KeyData",
          "inputValue": "/Ronne/key-data/cdw" },
        "Color Variable": {
          "inputType": "ScalarVariable",
          "parameterName": "Color",
          "inputGenre": "Variable",
          "inputValue": "/Ronne/scalar-vars/temp" },
        "Colormap": {
          "inputType": "ColormapVisAsset",
          "parameterName": "Color",
          "inputGenre": "VisAsset",
          "inputValue": "f52dabab-5b11" },
        "Glyph": {
          "inputType": "GlyphVisAsset",
          "parameterName": "Glyph",
          "inputGenre": "VisAsset",
          "inputValue": "1af03108-f1ed" } } } ] }
```

Listing 4.1: Data impressions of visualization state depicted by Figure 4.9. *UUIDs have been shortened to fit on the page.*

chapter five

Multi-Touch Querying on Data Physicalizations in Immersive AR

5.1 Introduction 61
5.2 Related Work 64
5.2.1 Representing Data in the Physical World 64
5.2.2 Exploring Data in the Virtual World 65
5.2.3 Early Tangible Interactions with Visualizations 65
5.2.4 Hybrid Virtual + Physical Systems 66
5.3 Hardware and Strategies for Addressing Technical Challenges . 69
5.3.1 Continuous Real-Time Registration of Virtual and Physical Displays 69
5.3.2 Sensing Multi-Touch Input on 3D Physicalizations 71
5.4 System and Interaction Techniques for Multi-Touch Querying on Data Physicalizations 72
5.4.1 System Overview 73
5.4.2 Streamline Seeding Interaction 74
5.4.3 Cutting Plane Interactions 74
5.4.4 World-in-Miniature Widget 74
5.5 Evaluation and Application 75
5.5.1 Characterizing Touch Accuracy 75
5.5.2 Antarctic Climate Science: Application and Feedback 77
5.6 Discussion 80
5.6.1 Limitations 80
5.6.2 Potential for Other Applications 81
5.6.3 Further Touch Sensing Evaluations 81
5.7 Conclusion 82

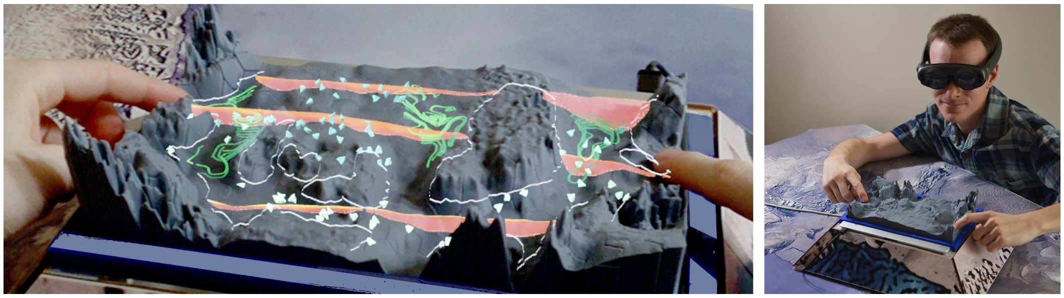

Figure 5.1: Adding direct multi-touch input sensing to a 3D physical printout and combining this with a perspective-tracked, stereoscopic AR display creates a hybrid visualization system that benefits from both the high-resolution, passive haptic feedback of the physical data printout and the interactive data exploration made possible by digital displays. Here, multi-touch queries create streamlines and cutting planes to explore a supercomputer simulation of ocean currents and ice melting under Antarctica's Filchner-Ronne ice shelf.

5.1 Introduction

[1]So far in this dissertation, we've established the utility of using physical rendering processes to encode data to enable better comprehension of and engagement with spatial data. The data physicalizations in Chapter 3 used a physical rendering process to translate digital data into physical data objects, while Chapter 4 introduced an "inverse" approach to the physical rendering process in which handcrafted, physical media created by artists are translated into elements of a digital extended reality visualization. We identify three key observations from these chapters that motivate further work in extended reality visualization and data physicalization.

Firstly, though the term "data physicalization" is relatively new, material properties such as size, form, and surface texture have been used throughout history for communication and education, including in the current digital age [Jansen et al., 2015]. Most modern data physicalizations leverage our natural predisposition for tactile data exploration like touching, holding, and rotating 3D data representations. The physical nature of the object can also enable better comprehension of the data by affording direct-touch interactions like marking points with a finger to compare values and judging distances based on a physical reference frame [Jansen et al., 2013]. So, based on the historic precedent for physical data representations and the contemporary evaluations that back up this precedent, we observe that data physicalizations are worth considering during a visualization design process.

[1] This chapter is based on work published at the 2021 ACM ISS conference under the same name. This work was accomplished in collaboration with Maxwell Omdal, Seth Johnson, Stephanie Zeller, Clara A. Richter, Francesca Samsel, Greg Abram, and Daniel F. Keefe [Herman et al., 2021].

Secondly, we observe that many extended reality technologies exist that attempt to replicate the embodied interactions we experience with physical objects in the real world; for example head-tracked, stereoscopic VR coupled with a glove robotic input/output device to provide touch feedback at the fingertips [Fang et al., 2015], or a VR display coupled with force feedback on the tip of a stylus [Jackson et al., 2012a]. Yet, for all the excitement and novelty currently surrounding digital extended reality technologies, they are still incapable of effectively reproducing the haptic benefits of running multiple fingers over the surface of a data physicalization such as using a finger to mark a point while comparing other areas of interest, or to follow a path on the physical surface [Jansen et al., 2013]. Thus, we observe that data physicalizations fill a gap in current visualization technologies that cannot yet fully engage the sense of touch.

Lastly, despite numerous successful examples of data physicalization, we observe what may be a fundamental limitation with respect to the requirements of scientific visualizations. Recall from Chapter 1 that scientists using these visualizations frequently need to make queries into the 3D multivariate spatial data in near real-time, and the visualization must provide rapid feedback which the scientist can use to explore, to test theories, and ultimately better comprehend their data. While recent physicalization work clearly demonstrates the ability to convey data from a variety of sources including melting glaciers [Segal, 2015], the predicted impact of future climate scenarios on native bird populations [Keefe et al., 2022a], global terrorism statistics [Madsen, 2018], and biological interactions of marine ecosystems [Miebach, 2008], none of these examples involve interactive data exploration. This raises the concern: Does the typically static, pre-fabricated property of physicalizations preclude their use for interactive tasks such as data querying, filtering, and calling up details-on-demand?

Researchers have explored two main approaches to overcome this seemingly inherent limitation. The first is mechanical: including a system of sensors to detect touch and linear actuators to physically move the 3D "pixels" of a display surface up and down in real-time [Leithinger et al., 2011]. These impressive actuated displays may indeed represent the future of physicalization work but presently remain costly and challenging to construct and transport. The second main approach is to create a hybrid display, combining physical elements with virtual elements to provide both interactivity and dynamic content (e.g., projection [Tateosian et al., 2010]). For this chapter, we define the scope of *hybrid visualization* [Jansen et al., 2013] (also known as *augmented physicalization* [Djavaherpour et al., 2021]) to include any visualization that contains physical components but adds interactivity using digital input and output mechanisms.

Building on this work, in this chapter we introduce a novel style of hybrid virtual + physical visualization designed specifically to support interactive data exploration tasks. We target the use case of multivariate, time-varying spatial data because this is a long-standing challenge for exploratory visualization research. This is one area where physicalizations have yet to be widely adopted, but have the strong potential to be beneficial because they make it possible to perceive data via multiple sensory channels simultaneously [Djavaherpour et al., 2021].

Our high-level goal is to combine the complementary strengths of digital visualization and physicalization while overcoming the individual limitations of each. Since physicalizations are meant to be touched with the fingers, we aim to support a level of interactivity that is similar to a digital multi-touch screen. Our first technical requirement is, therefore, to sense and respond to simultaneous touch input with multiple fingers placed directly on the surface of the physicalization. Further, the accuracy and responsiveness of this sensing should be similar to that of now-ubiquitous, multi-touch displays. However, unlike those displays, this input device should provide an added benefit of passive haptic feedback during interaction. Our second technical requirement is to robustly modify the visual display in response to the user input. Since we aim to support users analyzing complex spatial relationships in their data, meeting this requirement necessitates high-accuracy and low-latency alignment between the physicalization and the virtual data display so that users can make spatial judgements between the two (e.g., relating a flow feature displayed virtually to ridges in the bathymetry data displayed physically).

This research is significant because whether on planar surfaces, in immersive environments, or even hybrid virtual + physical systems, interactive techniques are essential to supporting exploratory visualization tasks, and so far these have been missing from data physicalizations. For 3D and 4D data, such traditional interactive techniques include user interfaces for 3D cutting planes, particle emitters, streamline seeding widgets, and multi-scale views [Nielson et al., 1997]. While implementations of these techniques exist for planar-display desktop setups (e.g., [Ahrens et al., 2005, Childs et al., 2012]), projection-based hybrid virtual + physical visualizations (e.g., [Piper et al., 2002, Alonso et al., 2018, Woods et al., 2016]) and immersive environments, including via multi-touch input (e.g. [Bryson, 1996, Coffey et al., 2011b, LaViola and Zeleznik, 1999]), our solution is the first to enable these interactive data exploration techniques in an immersive hybrid virtual + physical visualization.

Thus, this chapter contributes a novel approach to virtual + physical, hybrid, interactive visualization based on combining perspective-tracked stereoscopic Augmented Reality (i.e., immersive AR) with a creative approach to sensing touch *through* 3D printed data physicalizations. This work represents the first example of such a hybrid visualization that supports exploratory visualization tasks using classic interactive widgets for spatial visualization, such as particle emitters, streamline seeding widgets, and multi-scale views. The specific design developed (e.g., the custom tracking base) is justified via a combination of discussions of related work and results from formative experiments in the lab. The results are applied to create an exploratory visualization of Antarctic ice melt simulations actively studied by climate science collaborators (see Figure 5.1) who provide early feedback on the utility of the tool. This evaluation is complemented by a lab user study to characterize the lower-level accuracy of the physicalization touch input pipeline.

5.2 Related Work

In addition to the related work on data physicalizations and Tangible and Embodied Interactions discussed in Chapter 2, we discuss three further areas that contribute to our work in hybrid virtual + physical visualization systems.

5.2.1 Representing Data in the Physical World

Several advantages of data physicalization as compared to digital visualization have been demonstrated in the literature, including more efficient comprehension of data values when compared to 2D and 3D stereoscopic virtual displays of information [Jansen et al., 2013] and preference for working with physical visualizations over immersive VR visualizations [Danyluk et al., 2020]. Many physicalizations are also works of art; though their primary goal is often to engage and inspire, these artworks can also improve data understanding relative to more traditional visualizations [Djavaherpour et al., 2021]. Examples of physicalizations that are both artistic and improve data understanding range from 3D printing of volumetric datasets [Bader et al., 2018] to massive tripod data sculptures depicting climate change effects on native bird populations [Keefe et al., 2022a]. Researchers and artists believe that these may provide audiences with increased data engagement, with the added potential of emotional connection that may be an even more valuable outcome than traditional metrics of speed and accuracy in data comprehension [Wang et al., 2019].

Other data physicalizations exist solely for the pragmatic purposes of data analysis and interpretation. Such physicalizations have been applied to biological and medical data. Examples of this approach include the physicalization of molecular models [Carroll and Blauch, 2017] and the 3D printing of multi-color surfaces for analyzing the multivariate properties of pain phenotypes [Thrun et al., 2016]. Geospatial data are also commonly encoded with physicalization, with works ranging from tactile maps [Rase, 2011] to strategies for 3D printing an approximation of the Earth by dividing it into interchangeable, hierarchical cells [Djavaherpour et al., 2021]. While there are some indications that physicalizations are helpful for these pragmatic, exploratory visualizations (e.g., [Jansen et al., 2013, Danyluk et al., 2020]), there is also evidence to the contrary (e.g., [Berard and Louis, 2017]). In fact, Weiß et al. argue that there may be no significant difference in user engagement between web-based, VR, and hybrid physical visualizations for certain tasks [Weiß et al., 2020]. Although additional research is needed to continue to clarify the possible benefits, risks, and tradeoffs with data physicalization, it is quite clear that physicalization is not the right fit for every data analysis task; this further motivates our desire to explore the potential of hybrid visualization systems.

5.2.2 Exploring Data in the Virtual World

There is a rich history of research on interactive techniques for exploring spatial data using spatial and surface-based displays. Immersive data querying interfaces include flow visualization widgets, such as direct-manipulation 3D placement of particle emitters and streamline rakes [Bryson, 1996, LaViola and Zeleznik, 1999], interactive world-in-miniature widgets for improved navigation and manipulation [Acevedo et al., 2001, Coffey et al., 2011a], bimanual and multi-touch interfaces for slicing plane manipulation [Coffey et al., 2012, Schulze-Döbold et al., 2001], freehand and haptic-aided 3D sketching for data selection [Sobel et al., 2004, Keefe et al., 2008b, Jackson et al., 2012a], multi-touch convex hull selection on slicing planes [Coffey et al., 2011a], mid-air volumetric selection of sub-volumes of interest [Johnson et al., 2019a], and many more.

Many of these techniques translate directly from VR environments to AR, and both styles of immersive visualization have a history of successful application to geographic or topographic data (e.g., [Hedley et al., 2002, Bobrich and Otto, 2002]), as in our application, which combines flow and terrain data. Our approach therefore adopts many of the lessons learned from prior work, such as the importance of situating data queries relative to features in the spatial terrain and supporting the most common immersive flow visualization widgets. Of course, even the best of these purely virtual immersive environments lacks the real-world tangibility provided by data physicalizations [Jansen et al., 2015].

5.2.3 Early Tangible Interactions with Visualizations

Many visualization systems have built upon early research in Tangible User Interfaces (TUIs) (e.g., [Ishii and Ullmer, 1997]) to make the data interactions described in the previous section even more natural or intuitive [Shaer and Hornecker, 2010]. Our work shares a clear motivation with TUIs to bridge the gap between digital and physical elements. However, the "tangibles" in TUIs are typically different from the data physicalizations studied in this chapter. Most TUIs use comparatively simple tangible objects that act as proxies (e.g., pucks, tiles, bricks, pointers, or a doll's head). These have regularly been used to represent or select data; for example, using physical pointers to select virtual points on a map with an immersive AR overlay [Bobrich and Otto, 2002], exploring neural fiber data using a rolled-up piece of paper and stereoscopic panel display [Jackson et al., 2013], and performing query tasks in AR on 3D scatter plots with direct manipulation of physical props [Bach et al., 2018]. However, with a few noted exceptions (e.g. [Taher et al., 2015, Follmer et al., 2013]), users are not intended to perceive more than the position or orientation of the data by touching or viewing the tangibles. With data physicalizations, the role of the physical component shifts significantly to prioritize reading data, often using both the haptic and visual senses.

TUIs often integrate well with AR and VR displays [Billinghurst et al., 2008], and the technologies used to sense how the tangibles are touched and manipulated are also relevant to our work. For detecting touch with TUIs, there are two primary options: optical and direct. The optical approach, using cameras to detect touch contacts on the surface of a tangible, suffers from a well-known limitation that the precise moment of contact is difficult to detect [Wilson, 2010]. The precision of data exploration that is possible via this strategy is therefore limited, which is problematic when working with complex physical surfaces like geographic and bathymetric data. More direct touch sensing technologies (e.g., capacitive [Savage et al., 2012] and pneumatic [Tejada et al., 2020]) are therefore preferred from a standpoint of accuracy, but these more accurate sensing technologies are often difficult to integrate with digital fabrication techniques. For example, although flexible multi-touch sensors can be rolled to fit onto the surface of cylindrical tangible prop [Song et al., 2011], they cannot currently be stretched to fit over a more complex surface, even a sphere. We seek a solution that is direct-touch and high-accuracy but still integrates well with digital fabrication techniques.

5.2.4 Hybrid Virtual + Physical Systems

Moving beyond simpler tangibles to physical models that users touch and observe to read data values, the work that most closely relates to our own is in *augmenting* physicalizations with electromechanical or virtual components, an active topic of research also known as hybrid virtual + physical visualization. There are four categories of augmented physicalizations: surface augmentation, active physicalization, six-degree-of-freedom (6DoF) augmentation with a 2D display and 6DoF augmentation with an immersive 3D display [Djavaherpour et al., 2021].

In the first category, prior work using surface augmentation includes combining physical landscapes with GIS software [Tateosian et al., 2010], a table-scale physicalization paired with a digital energy data visualization [Kirshenbaum et al., 2020], and area-preserving displays for geovisualization [Dadkhahfard et al., 2018]. These surface-augmented physicalizations are often accomplished via projection (e.g. [Tang et al., , Tateosian et al., 2010]), though other technologies have been utilized (e.g., embedded LEDs [Taher et al., 2015]). Real-time registration between virtual and physical spaces is not usually required; often a single calibration is sufficient each time the system is assembled [Tateosian et al., 2010]. However, projection-based augmented physicalizations sometimes require projection mapping to ensure that no distortion occurs on the digital image – it has been reported as necessary in physicalizations over 6cm in height [Djavaherpour et al., 2021]. Surface augmentation is not ideal for our applications since we require understanding volumetric data, such as ocean currents at multiple depths.

Systems in the second category, active physicalizations, tend to require specialized hardware, and are usually limited to 2.5D applications where the mechanical elements are low resolution and move along a single axis. For example, Tangible CityScape [Tang et al.,] allows users to visualize multiple variables in the context of an urban environment physicalized with actuated columns, while EMERGE [Taher et al., 2015] enables visualization tasks like annotation, filtering, and navigation on 3D bar charts with actuated columns and color-changing LEDs. This approach is powerful and will perhaps one day evolve to create an ultimate display for surface-based 3D data. However, the current technology is limited in its ability to represent fine detail and in the prohibitive complexity and cost required to implement such displays. These factors along with a desire to complement physicalized data with spatially registered mid-air data steer us toward a hybrid solution using some form of AR.

In the third category, mobile phones, tablets, and desktop systems have been used to create a graphics viewed on the device's screen that overlay virtual content with a view of the physicalization captured by the device's camera. Examples include the tablet-based PLANWELL [Nittala et al., 2015] and desktop-based AR systems for analyzing structural molecular biology [Gillet et al., 2004, Gillet et al., 2005]. Unlike projection-augmented physicalizations, but similar to our work, these planar-display hybrid visualizations require real-time registration between virtual and physical spaces to allow a user to manipulate both the 3D viewpoint and the physicalization itself. In these situations, inside-out image-based tracking is often used (e.g. [Nittala et al., 2015, Gillet et al., 2005]), though other approaches exist (e.g. optical outside-in tracking [Pustka et al., 2012, Berard and Louis, 2017]). Techniques in this category can take advantage of the robust multi-touch input provided by the mobile device to support a variety of data interactions; however, they are performed on the mobile device's screen *not* on the physicalization itself, so the

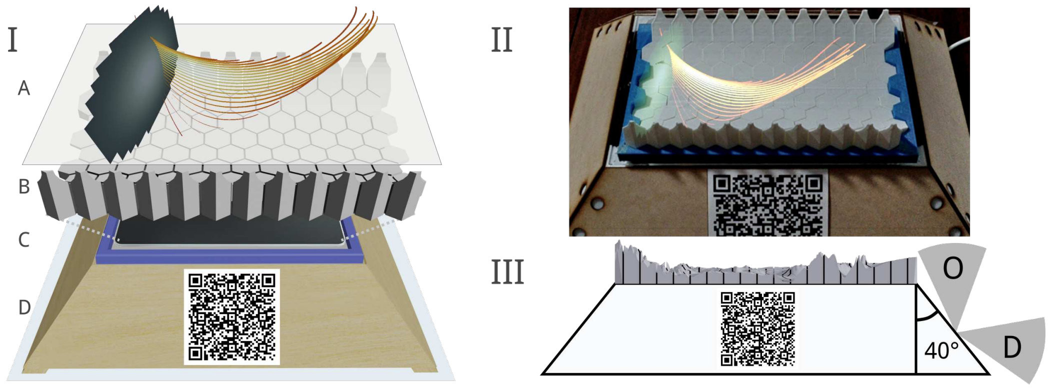

Figure 5.2: (I) An augmented reality visualization (A) is superimposed on top of a hexagonally segmented data physicalization (B). Force from a user's touches is transferred through the hexagonal columns to the force-sensing surface (C). The virtual visualization is spatially co-located with the physicalization using an image-based tracking system that determines 3D positions and rotations of image targets (D). (II) An example hybrid visualization containing a hexagonally segmented physicalization of a pipe and immersive augmented reality content showing user-chosen streamlines seeded based on a generated velocity field and a cutting plane selected with multi-touch querying. (III) The side angles of the truncated pyramidal base are angled at 40° to optimize tracking accuracy when viewing the physicalization from the side-on details view (D) and the top-down overview (O). Adapted from [Abawi et al., 2004].

potential passive haptic benefits of data physicalization are lost.

The fourth category of augmented physicalizations pairs 6DoF-tracked physical content with immersive, head-tracked stereoscopic AR virtual content. Prior work includes Augmented Maps [Bobrich and Otto, 2002] and a hybrid visualization of flood dynamics using 3D-printed topography [Zhang et al., 2020]. Systems in this style require robust, high-accuracy and low-latency 6DoF tracking for both the physicalization and the immersive AR device so that the virtual visualization remains spatially registered to the physical reference frame. We take this as a key requirement for our work and build upon the prior uses of image-based tracking methods (e.g., [Bach et al., 2018]). Although some prior work exists in this category, none meets our goal and requirement of supporting interaction through direct touch on the physicalization at a level similar to what can be provided by a typical multi-touch display.

5.3 Hardware and Strategies for Addressing Technical Challenges

Figure 5.2 presents an overview of our technical approach. By placing the physicalization on a custom base, we minimize tracking latency and maximize accuracy. By dividing the physical print into a regular grid of pieces that can move independently in the vertical direction, a pressure-sensing pad placed under the physicalization is able to sense touch events *through* the physicalization.

5.3.1 Continuous Real-Time Registration of Virtual and Physical Displays

Our strategy for consistently registering the position and rotation of the data physicalization relative to the user's eye uses a passive inside-out optical image tracking system. Although this is the best match with the Magic Leap AR hardware we used, other 6 DoF tracking solutions were considered, including electromagnetic, passive outside-in, and inside-out 3D model (as opposed to image) tracking. Passive inside-out tracking based on planar image targets is the most portable, only requiring a single camera for capturing the scene and images affixed to each object to be tracked. Its sole requirement is a target image placed near the physicalization such that it is in the line of sight of the camera simultaneously with the physicalization, as in the image target in Figure 5.2(I-D). Coverage and portability of inside-out image-based optical tracking also surpasses that of its outside-in counterparts because it is not bounded by a "tracking volume" – the space where all external cameras have a line of sight – and instead, the images and cameras are connected directly to the physicalization and the viewer, and thus will always be available without setup. This also means that it is not necessary to track both the user's head position and the physicalization; the relative offset from the image target to the tracking camera is sufficient to place the physicalization in 3D space. It is possible to perform an initial registration with image targets, then switch to the AR HMD's onboard 6DoF tracking to maintain the alignment. The onboard tracking is lower-latency, but it does not allow a user to manipulate the physicalization interactively, for example, to spin it around to view the other side. In practice, we found that the Magic Leap's internal room-scale tracking system was not accurate enough to maintain sufficient alignment between the virtual and physical content of the visualization without continuously enabled image tracking.

The key to optimizing an image-based tracking system is to understand the tradeoffs in design choices for the number of images used to track objects and their placement in physical space. Thus, our design decisions focus on these two issues.

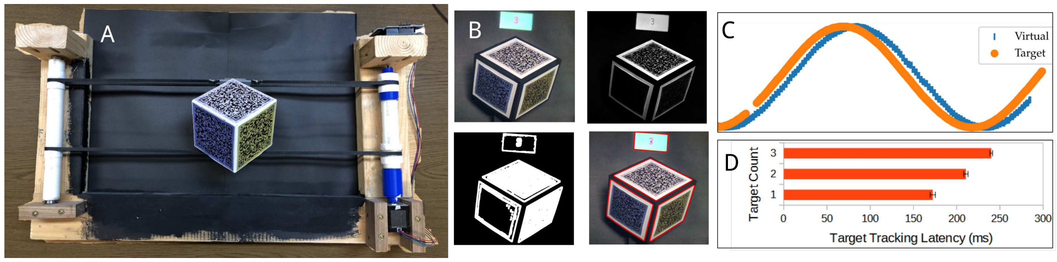

Figure 5.3: The experimental setup to optimize latency for image-based tracking includes (A), an oscillating image target, (B) boundary extraction for the virtual and physical object positions, and (C) positional samples of virtual and physical target used to determine latency. Image (D) shows the latency results based on number of targets.

To optimize the number of images used to track the physicalization, we conducted an experiment to measure latency of the tracking system based on the number of image targets in view of the tracking camera. Figure 5.3A shows the approximate view of the AR device facing the sinusoidally oscillating image tracking assembly with a variable number of image targets attached. Video was captured through the left eye of the AR device of a virtual object's position tracked with the physical object's location, and computer vision methods were used to extract the true positions of the physical and virtual objects in the scene (Figure 5.3B). Lastly, following Steed's Sine Fitting method, both the virtual and physical object trajectories were fit to a sine curve, and the latency of the system was calculated as the phase distance between the two curves [Steed, 2008]. We found that the latency of the image tracking system does indeed increase with the number of simultaneously visible targets; we found a latency of 172ms, 211ms, and 240ms for 1, 2, and 3 targets, respectively. This suggests the physicalization tracking will have the lowest latency when tracked by just a single image. Using positional accuracy measurement methods based on Groves et al. [Groves et al., 2019], we verified that tracking just a single image is sufficiently accurate, providing sub-millimeter positional accuracy for a comfortable viewing distance of 30cm. However, we found it impossible to find a position where a single image could be placed next to the physicalization and force-sensing tablet while also providing reliable tracking from all viewing angles.

This led us to consider a stage or stand for the physicalization that provides a good tracking target from any viewpoint but limits the view to one (ideal) or at most two image targets at a time. Our starting point was a box with one image on each vertical side. However, noting that prior work has shown that image targets track most accurately when viewed from non-perpendicular angles [Abawi et al., 2004], we further refined the design to the truncated pyramid shape pictured in Figure 5.2. The four-sided, truncated-pyramid base minimizes the number of targets simultaneously visible from any viewpoint, while tilting the sides of the base to an angle of 40° from vertical, which, based on the previously reported data [Abawi et al., 2004], should optimize tracking for the top-down and side-on perspectives, as illustrated in sub-figure (III).

Our implementation uses commercially available software and hardware. A Magic Leap is used to display the augmented reality content with a resolution of 1280 x 960 pixels per eye, a horizontal field of view of 50°, and a 120Hz refresh rate. Optical tracking uses Vuforia [PTC, 2020]. Graphics are rendered using the Unity game engine, and a custom shader is used to occlude the augmented reality graphics with the physical model, as shown in Figure 5.1 *right* (though a user's fingers do not occlude the graphics). It should be noted that the images in this chapter containing AR visualizations were generated using the Magic Leap's built-in screen capture system, which had a screen space offset between real-world and virtual content of approximately -0.037 units in the x direction and -0.006 units in the y direction. This was compensated for by visually calibrating the scene to a physical target before making the screen captures.

5.3.2 Sensing Multi-Touch Input on 3D Physicalizations

Our design for sensing touch through a 3D printed terrain-like surface builds upon related work and workshops on the theme of "Touching the 3rd Dimension" [Steinicke et al., 2012]. Extending the fluidity of multi-touch interaction off the surface of the screen and into/onto spatial displays is a longstanding goal of this research, which includes work on perceptual illusion [Bruder et al., 2011], tracking 3D gestures anchored to a surface [Jackson et al., 2012c], redirected touching in purely virtual environments [Kohli, 2010], and even dynamically inflatable multi-touch surfaces [Stevenson et al., 2010].

Figure 5.2 shows our solution for translating physical touch into virtual environments. To enable multi-touch interactions directly on the physicalization, it is split into a series of hexagonal columns with a radius of 0.75cm. This size of hexagon is chosen to be small enough to approximate the radius of a human fingertip yet maintain a manageable number of columns for the entire physicalization. Tiled hexagons are used to avoid the rigid visual structure that a regular square or rectangular grid imposes on the physicalization.

It should be noted that while the resolution of each column is fairly coarse, sub-hexagon touch precision is achieved by placing the partitioned physicalization atop a high-resolution force-sensitive surface, a Sensel Morph [Sensel, Inc.,]. The hexagonal columns have flat bottoms which enables our software to interpolate a user's touch position within a hexagon. The Sensel Morph has a size of 24cm x 13.85cm and contains 185 x 105 force sensing elements which detect forces 0.049N-49N.

Using polygonal columns on such a sensitive device presents its own series of challenges. Firstly, as the user presses down on a particular column, multiple touch points

Figure 5.4: (Left image) The physicalization and force-sensing surface are placed on a truncated pyramidal base. The base is situated onto a large fabric poster with contextual NASA satellite imagery around the Filchner-Ronne ice shelf in the Antarctic. *The imagery is not to scale with the physicalization.* (Right image) Worlds-in-miniature allow the user to view their data at multiple scales. All interactions still take place on the physicalization; when users are focused on the larger world, the physicalization acts as a haptic proxy.

may be registered by the device from the column's edges contacting the surface. In order to remove these "ghost" touches and ensure there is only one touch point registered for each finger, configurable parameters are used. Since the physicalization is divided into 0.75cm columns, an assumption is made that two fingers will never touch the same column – if a cluster of touch points within a diameter of 1.5cm is registered, the point inside this cluster with the greatest force is kept and the remainder are culled. To remove any other ghost touches remaining from the weight of the physicalization or other physical factors, the touch points are thresholded such that those with at least 5x less force than the others and those with an absolute force of less than 0.98N are ignored. Once the touch points are filtered, they are made available on a network using VRPN [Taylor et al., 2001] in order to be accessed by other devices such as the augmented reality display.

In addition to the well-developed application to Antarctic ocean current simulations described in detail later, we have also successfully applied these strategies for multi-touch input and tracking to other datasets, such as the flow field in a pipe, pictured in Figure 5.2.

5.4 System and Interaction Techniques for Multi-Touch Querying on Data Physicalizations

This section describes how three classic interactive data querying techniques can be re-implemented in the new context of hybrid virtual+physical visualization. Before describing these in detail, we begin with a brief overview of the system architecture

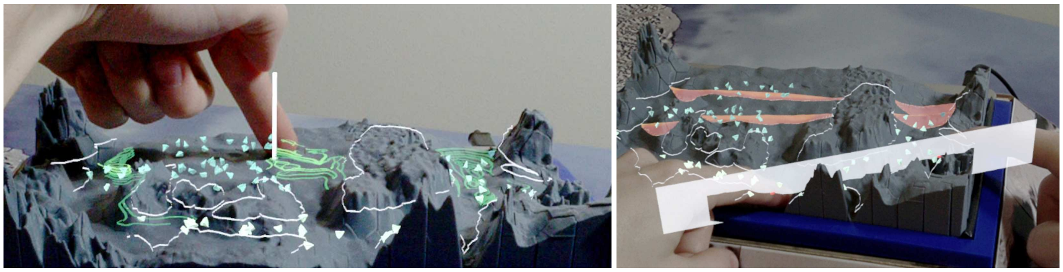

Figure 5.5: Streamlines are interactively seeded by moving a finger around on the surface of the physicalization and releasing (left), while cutting planes are created by using two fingers (right).

used in our implementation.

5.4.1 System Overview

Our system architecture includes four components: the immersive augmented reality visualization, the multi-touch physicalization attached to a computer, the data host, and the visualization host.

Let us motivate this multi-component system by following the steps a user might take when making queries into a visualization. The user begins by donning an AR device in a space near the data physicalization. Figure 5.1 shows an approximate view of what the user might see here; the immersive AR content for that visualization is spatially registered to the physicalization using optical image-based tracking of the target images shown in Figure 5.4 (left). When the user touches the physicalization, this input is sensed as described in Section 5.3.2 then relayed to the AR device by a VRPN server on the touch host. The AR headset in turn calculates the proper virtual position of each touch point. After converting the touch points from room-space to data-space using a series of affine transformation matrices [Jackson and Keefe, 2019], a VTK data query is synthesized and sent to the data host as an HTTP request. Query results are calculated on the data host using VTK [Schroeder et al., 2004], and a geometric representation of the data is sent back to the AR headset and displayed in the visualization.

While not the focus of our human-computer interaction research, a companion touch screen interface is also supported in this architecture for input related to system control (e.g., specifying a timestep to view or switching between query modes). Additionally, we provide a touch-enabled design user interface that allows all data-to-visual mappings on the AR visualization to be modified interactively [Herman et al., 2020]. The AR visualization is rendered using the Artifact-Based Rendering (ABR) technique [Johnson et al., 2019b], delivering visually rich graphics to complement the data physicalization.

5.4.2 Streamline Seeding Interaction

Figure 5.5 (left) depicts a demonstration of an interactive stream rake widget that seeds streamlines along a line segment defined by a point on the surface of the physicalization and a second point directly above it. Streamlines are one of the essential tools for exploring fluid flow data from computational fluid dynamics and other numerical simulations [Nielson et al., 1997]. In particular, stream rakes are used to create and view a collection of flow lines that pass through multiple seed points that are evenly spaced along a user-defined line segment. Variation between the resulting flow lines can then show researchers how flow patterns change along the axis defined by the line (ocean depth in our case).

In our implementation, the base of the rake is defined by a user's touch point on the physicalization and the up vector. As the user moves their finger along the physicalization, a preview of the seeding line is displayed on the AR device and updates in real time. The streamline resolution, number of seed points, and length are all adjustable parameters. When the streamline rake is selected, its coordinates are converted into data-space. Coordinates are sent to the data server, where VTK components are used to advect the streamline geometry. The streamlines are then sent back to the immersive AR visualization application, as shown in Figure 5.5 (left).

5.4.3 Cutting Plane Interactions

Figure 5.5 (right) demonstrates a cutting plane query defined by a user's two touch points on the physicalization. Interactively creating cutting planes is another useful tool for exploring 3D volumetric data [Klein et al., 2012]; they allow a user interactively define a 2D slice into the data.

In our system, the user's two touch points and the up vector define a basis for the cutting plane. Again, an interactive preview of the plane is updated in real-time as the user moves their fingers. When these touches are released, the coordinates are converted into data-space and sent to the data server. VTK components are used to create the cutting plane geometry which is then returned to the immersive AR app.

5.4.4 World-in-Miniature Widget

The ability to explore data at multiple scales is another indispensable technique in scientific visualization [Nielson et al., 1997]. Worlds-in-miniature [Stoakley et al., 1995] is a technique that makes it possible to view a virtual environment at different scales while simultaneously interacting with a proxy that is at a comfortable scale. This technique has been used extensively in visualization research, for instance to compare forest plots

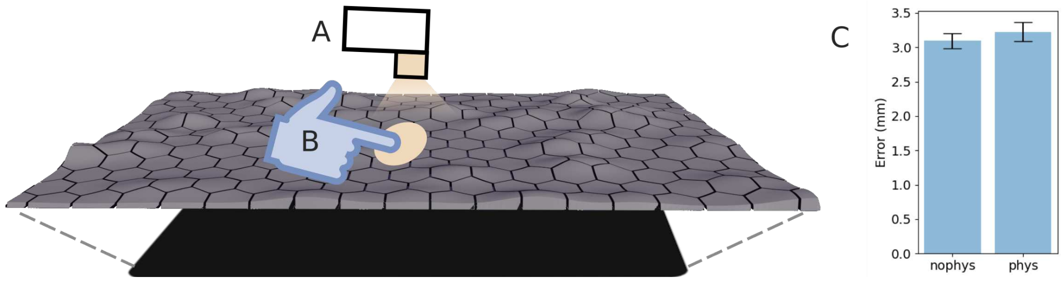

Figure 5.6: Touch accuracy was measured with and without the tiled physicalization present. Random Cartesian coordinates are projected on the bare force-sensing surface or physicalization (A), and a participant touches these points (B). Image (C) shows positional error across conditions. Error bars show CI95.

over time [Nam et al., 2019] and to explore flow data via multimodal displays [Coffey et al., 2011a].

Figure 5.4 (right) shows our implementation of worlds-in-miniature using the hybrid virtual + physical visualization model. The hybrid worlds-in-miniature technique uses the data physicalization as both a primary view for the data and as a haptic proxy [Nilsson et al., 2020] for a larger AR visualization.

5.5 Evaluation and Application

Returning to the goals and requirements stated earlier, supporting exploratory visualization tasks requires a robust, high-accuracy pipeline for capturing multi-touch input and a responsive display that can usefully update the digital aspects of the visualization while maintaining accurate registration of the digital and virtual components. This section reports on two complementary evaluations: 1. a quantitative lab study on touch accuracy, and 2. an application to Antarctic climate science data and feedback from a domain scientist.

5.5.1 Characterizing Touch Accuracy

We conducted a user performance study to better understand the low-level capability of sensing multi-touch input *through* a data physicalization. During the touch study, the AR device was not included because the touch coordinates reported to the Touch Host and AR device are identical. Touch data were collected using the experimental setup pictured in Figure 5.6, which was designed to detect touches at specific (x,y) coordinates with and without the hexagonally tiled physicalization present. The study was designed as a within-subjects, repeated measures A vs. B comparison of two conditions: touch input sensed through a hexagonally-tiled physicalization (*phys*) and touch input sensed directly on the Sensel Morph's touch pad (*nophys*). From our experience, we hypothesized that the *nophys* condition would be close, but slightly more accurate than the *phys* condition.

Methodology

Participants ($n = 8$) were recruited from the local computer science student population. All participants were right-handed and all reported having normal or corrected-to-normal vision. After informed consent, the participants completed a series of pointing tasks in which they were instructed to touch a target as accurately as possible with their dominant-hand index finger. The tasks were presented in three blocks of 30 trials each. The first block was used for training and consisted of 30 trials in conditions alternating between *phys* and *nophys*. During the second and third blocks, participants touched points on the bare force-sensing tablet (condition *nophys*) and the physicalization atop the tablet (condition *phys*). The order of these two blocks was counter-balanced across subjects to ensure minimal learning effects.

Target touch points were generated using a random uniform distribution across the 24cm x 13.85cm physical surface. Each random point was then displayed sequentially to the participant via a projector aligned with the force-sensing surface (see Figure 5.6 A); each point was revealed after the participant had completed the touch action for the prior point, and points were visualized with a diameter of 5mm. The participant's touch points were recorded after filtering to remove ghost touches (using the strategy described in Section 5.3), and additionally adjusting for the Sensel Morph's hardware offset (x: 3.2mm, y: -2.6mm) and scale factor of 0.97.

The physicalization used in the study was of a synthetic terrain dataset and was the same for all participants. The terrain data were generated by adding 50 randomly distributed Gaussian peaks to create a bumpy surface with a maximum height of 8mm (Figure 5.6 B). These synthetic data were used, rather than the pipe or Antarctic data presented elsewhere in the chapter, because we determined in pilot testing that a flatter surface was needed to avoid distortion effects when projecting the target points on the physicalization.

Results

Touch data were recorded for *8 users* $\times$ *30 points*, totalling 240 data points for each condition. 18 duplicate touches (touches that registered within 0.5s of each other) were registered for *nophys* and 38 were registered for *phys*; as is common in touch-tracking software, these duplicates were averaged and counted as one touch. Additionally, 10 expected touches for *nophys* and 4 expected touches for *phys* were not detected as they were either within the dead zone surrounding the force-sensing tablet or below the force threshold in our touch reporting software.

Touch accuracy was measured by calculating the `touch_error` as the Euclidean distance between the ground truth Cartesian coordinates and the coordinates of a participant's touches. Four points with an x or y coordinate more than 3σ away from the mean were considered outliers and discarded.

As shown in Figure 5.6C, participants selected points more accurately under the *nophys* condition ($\mu = 3.09$mm, $\sigma = 1.64$mm) as compared to the *phys* condition ($\mu = 3.22$mm, $\sigma = 2.1$mm). A Shapiro-Wilk test showed that the data were not normally distributed ($p < 0.0001$). Thus, a Wilcoxon signed-rank was used to analyze significance. The two-sided Wilcoxon signed-rank test indicated that the physicalization did not have a significant effect on `touch_error` ($p = 0.753$); thus, we fail to reject the null hypothesis that there is no difference in mean `touch_error` between conditions.

Interpretation and Discussion

Comparing the results of `touch_error` between conditions indicates that the average Euclidean distance between a participant's touch points and the ground truth points was slightly greater in the *phys* (3.22mm) condition than in *nophys* (3.09mm). We did anticipate some difference in this direction since the rigidity of the physicalization hexagons certainly contributes some error. However, after observing participants during the study, we also suspect that some error was due to shadows in the projection used to display touch point as well as the physical pixel size (about 0.5mm x 0.5mm).

These combine to leave us surprised that the study was unable to detect a significant difference in touch accuracy across approximately 240 trials. Other factors could also impact this result. For example, the number of participants was small and their backgrounds rather homogeneous (limitations necessitated by the constraints of the COVID-19 pandemic).

Despite these factors, we take the study as a promising initial characterization of the feasibility of the core concept behind our tiled physicalizations–that the force of a participant's touch can be transferred directly through the column to the tablet, opening the door for sensing touch through the type of rigid 3D printed physicalizations that are most easily created today.

5.5.2 Antarctic Climate Science: Application and Feedback

To assess the potential of multi-touch querying on data physicalizations in a real-world scientific context, we constructed the augmented physicalization pictured in the examples throughout this chapter and conducted a user interview with a climate scientist

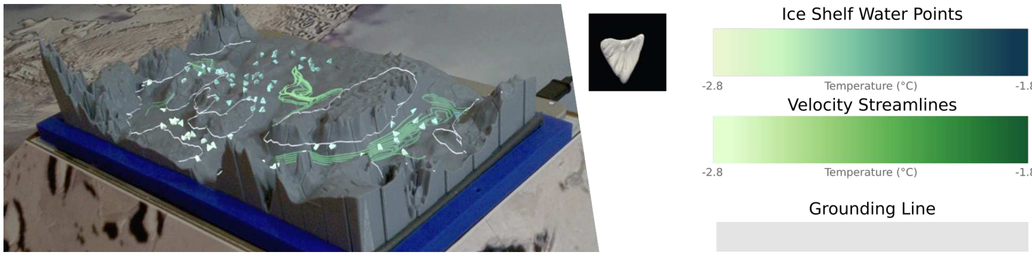

Figure 5.7: Streamlines are seeded based on the velocity of the ocean currents underneath the Filchner-Ronne ice shelf, and are colored by the ocean temperature. Glyphs are distributed based on a density-based point sampling of the Ice Shelf Water mass. The grounding line where the ice intersects with the water and land is shown as a ribbon on the visualization.

collaborator actively studying the phenomena. The Filchner-Ronne ice shelf in Antarctica is of significant interest to climate scientists because ice shelves extending from the Antarctic continent prevent ice sheets on the continent itself from melting rapidly and flowing into the sea (see Figure 5.7). The water content of the Antarctic ice sheets and shelves combined has the potential to inundate many of the planet's major coastal population centers if completely melted. The melt rates beneath the ice shelves exceed the surface melt rates and are difficult to monitor, as water masses and currents of varying temperature, salinity, flow rate, and more must be tracked and are invisible from the surface. A more detailed understanding of the currents, their intensity, and their temperature will enable scientists to more accurately predict the melt rate of the ice shelves, which play a critical role in sea-level rise as well as climate patterns across the globe [Petersen et al., 2019, Jeong et al., 2020, Abram et al., 2021].

Technical results of the application to climate science in Antarctica's Filchner-Ronne ice shelf are pictured throughout the chapter as well as in the accompanying video. All three of the classic data querying interaction techniques described in Section 5.4 are applicable to these data and are controllable via multi-touch queries made on top of the Filchner-Ronne ice shelf 3D print.

Data Conversion

Due to the physical constraints of the direct-touch sensing mechanism, a subsection of the Filchner-Ronne ice shelf was selected based on conversations with collaborating climate scientists about areas of interest within this region. The subsection measures 945 km by 545 km, and underwent transformations from its original coordinate space in the climate simulations, including a stereographic projection, a vertical exaggeration of 60x, clipping and scaling to the size of the force-sensing tablet, and finally segmenting into hexagonal columns as discussed in Section 5.3.2. Special care was taken to preserve the original directionality of the ocean current flow vectors when applying the stereographic projection and performing affine transformations on the data. Hexagonal columns of the surface data were 3D printed with a layer height of 0.15mm, taking a total print time of about two days for the 247 hexagons in the segmented surface data.

Figure 5.4 (left) shows the 3D printed region of interest along with NASA satellite imagery from the Filchner-Ronne ice shelf. Tracking images that align the virtual and physical visualizations are adapted from NASA satellite imagery, which provides additional context for the position and orientation of the virtual data and the physicalization relative to a bird's-eye view of Antarctica. The images are affixed to the truncated pyramidal tracking base, and the context they provide is enhanced by placing the entire assembly on top of a large 2D poster, extending the Antarctic imagery.

Climate Scientist Interview

The interview was conducted with one investigator physically present and two remote investigators joining via video conferencing. We began the interview process by asking the climate scientist to explore the system, using a "think-aloud" approach as we directed him through a series of introductory tasks designed to demonstrate the capabilities of our system. These tasks included seeding streamlines, creating cutting planes through the volumetric data, and using the multi-scale view to view details of the ocean current mixing. We then prompted the scientist throughout his exploration process with questions designed to encourage both constructive criticism about the system as a whole and feedback on the scientist's usage of the system. Questions included, for example, "How does it feel to explore your data in this format?" and "How do you think this configuration could play a role in your collaboration with peers?"

When setting up the physicalization for the interview, we noted that the orientation of the physical data object relative to the user is a critical consideration [Sauvé et al., 2020]. The bathymetric physicalization was oriented with North facing away from the user, a common perspective that the scientists use when working with these data.

The interview revealed several preliminary insights. Following the introduction tasks, the climate scientist conducted a rapid "lay-of-the-land" exploration by touching the printed bathymetry in several locations, seeding clusters of streamlines to understand the velocity, temperature, and direction of interesting flow patterns. He noted our system's utility in selecting streamlines precisely on the surface and added that automatically seeding an initial set of streamlines at the start would help narrow down areas of interest in the data.

The climate scientist reacted to the physicality of the data visualization, saying "*The fact that the data overlays something physically is really wonderful...I've never experi-*

enced this before. There's a really intuitive feel to it." and "*[The printed bathymetry] really grounds you and provides useful context for the data itself.*" He noted the system's potential for improving his group's research and inter-scientist communication more generally, particularly if a future iteration of the system acted as a centerpiece that team members could walk around, examine, and reference interactively on a tabletop. He further indicated that a hybrid system might increase public interest and improve communication of critical but complex concepts in this field.

In addition to these overarching comments, the climate scientist provided useful feedback on several specific aspects of the design. He observed of the hexagonally-segmented physicalization itself, "*I like the hexagons... it's much more stable [than squares], and I think having squares instead of hexagons would be much more distracting.*" The scientist expressed particular interest in the grounding line – the line where water, depressed by surface ice, meets the land. This feedback prompted the grounding line's inclusion as shown in Figure 5.7. He also stated that a larger physicalization of the bathymetry would have been extremely useful, even if it wasn't all touch-sensitive, adding that he found himself primarily using the printed bathymetry as an orientation device as opposed to searching for small details in the physicalization. This observation was striking. Counter to our expectations, the physicalization became a spatial marker in three dimensions that acted as a peripheral canvas for the data of greater interest as opposed to a primary data object itself.

Although the subjective user feedback presented here should be treated only as an early data point in understanding the potential of this approach, the interviewed scientist and other collaborators noted several use cases where there is strong potential for hybrid virtual + physical visualization. Speaking on the subject of collaboration, one scientist remarked "*Having it available to work with and show other scientists – I think there's a lot of potential there.*" Furthermore, "*There's a lot of potential for people to see and get excited about the science because it's something new and interesting.*"

5.6 Discussion

5.6.1 Limitations

Our approach to hybrid visualization, while a novel step forward, has two principal limitations. First, for some workflows, the digital fabrication time may be prohibitive. Printing time can be reduced by increasing the layer height, but this comes at the cost of a lower-resolution model. However, our scientist feedback interview indicates potential for some flexibility in resolution, depending on how a user identifies and works off of the physicalization.

Another issue that emerged with the hexagonal column design is the challenge of keeping all the tiles linked together. Researchers needed to be cautious when moving the physicalization from the 3D printer to the touch surface, as the hexagons could become a 3D puzzle very quickly if dropped. Currently, we mitigate this issue with the 3D printed frame shown in blue on Figure 5.2 I-C, which nicely holds the tiles in place on the surface but will not prevent them from scattering if turned upside down. Although several prototypes for linkage between hexagons together while preserving complete touch sensitivity have been made, we have not yet found a foolproof resolution.

5.6.2 Potential for Other Applications

Our solution to hybrid virtual + physical visualization has been applied to several datasets so far, including those shown in this chapter, though our solution is not limited to such data. We have considered, for example, applying this technique to brain microstructure data as well as treatment-planning data from radiation therapy procedures. In these anatomical cases, we would physicalize a slice or surface out of a CT scan and segment it, enabling the domain scientists to create cutting planes and interactively seed directional glyphs at points of interest to further explore their data.

The addition of a force-sensitive surface underneath the physicalization also opens up new possibilities for interactive modular techniques that not only leverage a user's *touch* but also the touch's *pressure*. This is not a new area of research [Buxton et al., 1985], but its application to 3D exploratory visualization could inspire new data querying widgets, such as sampling glyphs with density determined by pressure or using pressure to control the number and length of streamlines. Further, such systems present new and highly complex visual design challenges and open new avenues for experimentation, including the materiality, appearance, and scale of the physicalizations and their interactions with the virtual visualization.

5.6.3 Further Touch Sensing Evaluations

One of the most interesting discussion points raised by the touch accuracy study is the issue of what accurate touch means when sensing touch *through* a physicalization. The results from the study in Section 5.5 suggest that there are many contributing factors, both hardware and software. While the parameters of the touch-filtering algorithm remained constant throughout the study and application, these are all individually adjustable to achieve a more accurate touch and further avoid duplicate and missed touches. Such tuneable parameters may include the clustering diameter, minimum absolute force threshold, and force threshold relative to other touch points on the surface. Beyond these software parameters, we must also consider the hardware setup – what is the effect of polygon size on touch accuracy? Polygon shape? Height of the columns? The durability and linkage between columns? This last item proved to be of particular interest; during the touch study, several participants voiced concerns about "breaking" the physicalization during the *phys* condition, and one managed to unseat a hexagonal tile while touching the physicalization. These are all factors that have the potential to impact accuracy of a multi-touch physicalization, and thus need to be evaluated in future studies. Additional studies must also be run (e.g. in the style of [Jansen et al., 2013]) to test the exploratory capabilities and interactive techniques used in the hybrid visualization system.

5.7 Conclusion

In this chapter, we presented a new approach to multi-touch sensing on data physicalizations and combining physical and virtual data together into a single, interactive queryable visualization. The motivation for the work was to fuse the "best of" physicalizations and virtual displays; specifically our goal was to create a form of data physicalization that can be used for interactive, exploratory visualization. We demonstrated that this is indeed possible; in fact, we showed for the first time what three classic spatial data exploration widgets might look like when designed for use with a data physicalization in conjunction with immersive AR. The preliminary quantitative lab study of touch accuracy lends support to the feasibility of the core concept of sensing touch input through a tiled physicalization surface, and the application to exploratory visualization of Antarctic climate simulations demonstrates how the sensing and interaction techniques can be used with real data. We believe the sensing solutions and interactive techniques presented here are a major step forward for the future of hybrid visualization platforms. This exploratory research provides evidence for significant future potential of interactive, surface-based hybrid visualization. As a whole, this work provides the first example of such a physicalization that supports the style of back-and-forth human-in-the-loop needed by today's scientific sensemaking processes, and comments from the scientist feedback sessions showed that this visualization modality has a strong potential to make scientific data both more comprehensible and more embodied.

chapter six

Touching the Ground: Evaluating the Effectiveness of Data Physicalizations for Spatial Data Analysis Tasks

6.1 Introduction **84**
6.2 Related Work **86**
6.2.1 Evaluating Visualizations 87
6.2.2 Comparing Physicalizations and Digital Visualizations 89
6.3 Study Design **90**
6.3.1 Hypotheses 90
6.3.2 Modalities and Visualization Design 92
6.3.3 Data 95
6.3.4 Tasks 96
6.3.5 Experimental Procedure 99
6.4 Results **103**
6.4.1 Quantitative Measures 103
6.4.2 Qualitative Feedback 110
6.5 Discussion **112**
6.5.1 Interpretation of Results 112
6.5.2 Implications of Physicalization on Spatial Data Analysis . . . 115
6.5.3 General Discussion 118
6.6 Limitations and Future Work **119**
6.6.1 Limitations 119
6.6.2 Future Work 120
6.7 Conclusion **121**

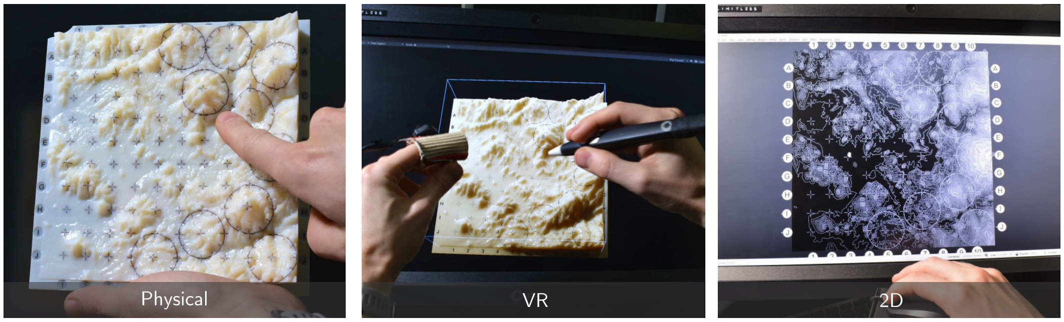

Figure 6.1: Data physicalizations provide many benefits over digital data displays, including haptic perception and providing an embodied reference frame for the data. This chapter compares the effectiveness of spatial data physicalizations (left) with typical virtual reality (middle) and 2D visualizations (right).

6.1 Introduction

For millennia, humans have used physical objects to help us make sense of the universe. Physical objects have played a vital role in new scientific discoveries and scientific sense-making processes throughout their history, from discovering the nonlinear relationships between volume, energy, and entropy in thermodynamics [Kriz, 2007] to breakthroughs in understanding the function of myoglobin storing oxygen within animal cells [Kendrew et al., 1958]. Data physicalizations convey data through their geometric and material properties [Jansen et al., 2015], and their capacity to support exploration and feedback through touch enables yet-unparalled perceptual advantages compared with today's most advanced digital displays. Namely, physicalizations afford high resolution physical rendering, suffer zero system latency, and enable physical interactions with data; they make it possible to directly touch one's data and make comparisons between points, to easily move and rotate a physical object, to engage multiple senses with the data, and ultimately to take advantage of our own embodied human reference frame to make judgments of size and distance [Hornecker, 2011].

There is a strong need to provide more evaluation-driven evidence to substantiate these *theorized* benefits of data physicalizations, both by creating new evaluation methodologies and by validating physicalizations against proven evaluation measures. Some of data physicalizations' benefits are well-documented through innovative means, including their memorability [Khot et al., 2014], users' engagement with the data [Holstius et al., 2004], or even the sense of awe and emotion that they inspire in viewers [Segal, 2015]. These physicalizations welcome a new era in visualization evaluation where the user experience, engagement, emotion, and embodiment of visualizations are valued alongside their comprehensibility with far-reaching results that speak to the humanity of data and visualization. Although physicalizations often push the boundaries of new and innovative types of evaluation, many researchers remain unconvinced of the value physicalizations provide due to the lack of "traditional," performance-driven evaluations in the physicalization community.

Typical visualization research performance evaluations measure the *time* and *errors* of participants completing specific tasks with different visualization techniques, and comparing quantitative results to determine which technique is better [Jackson et al., 2012b]. In fact, such performance measures have been the main subjects of numerous conceptual models highlighting the value of visualizations [van Wijk, 2005, Munzner, 2009, Stasko, 2014]. Evaluations in this traditional analytic style are valuable for determining the *efficiency* and *performance* of visualization systems; however, these are not the only metrics for the viability of visualizations [Wang et al., 2019]; indeed they rarely encompass all the components that make a visualization truly *effective*. Despite their limitations, time and errors performance evaluations are still the predominant measure with which different visualization techniques are compared, and very few instances exist of these evaluations on data physicalizations. Thus, visualization practitioners remain hesitant to adopt physicalization as a primary modality of data analysis because its performance advantages are not yet well-evaluated compared with current digital visualization technologies, and critics of data physicalization remain unconvinced that the additional time and material costs of producing a physicalization are worth it for the theorized benefits physicalizations provide.

As a first step to validate the theorized perceptual and embodied advantages of data physicalizations, a small number of prior studies have evaluated the performance of data physicalizations compared with digital visualization modalities. Jansen et al. found that participants took much less time solving tasks with physical 3D bar charts than with 3D desktop visualizations [Jansen et al., 2013], and Danyluk et al. found similar results with task time when comparing physical 3D bar charts with VR visualizations of the same charts [Danyluk et al., 2020]. While these evaluations show promising results for the utility and performance of using data physicalizations for data analysis tasks, further evaluation is required to validate and generalize their results beyond the specific data and tasks that were used.

Specifically, more evaluation is needed for physicalizations that target *spatial* data. While many examples from the literature physicalize spatial data, they are predominantly data sculptures intended for public consumption [Lozano-Hemmer, 2004, Meier, 2017] or proof-of-concept systems with no comparison between other visualization types [Gillet et al., 2004, Le Goc et al., 2019, Ang et al., 2019]. These types of physicalization are important and necessary for advancing the state of physicalization research and communicating data in new ways, but their existence does not negate the fact that there is an alarming lack of evaluation for spatial data physicalizations: out of 33 data physicalizations categorized as "SciVis" (spatial data) in a recent survey paper [Djavaherpour et al., 2021], only 10 contained a user study evaluation of any sort, and *not a single one evaluated spatial data physicalizations with respect to existing digital visualization techniques*. Many factors contribute to the current lack of evaluation in physicalization with spatial data, not the least of which being the inherent differences between spatial and non-spatial data. Spatial data are distinguished from non-spatial data by the fact that their structure (i.e., placement in space) is *given* by the data, as opposed to non-spatial data where the data's placement in space can be chosen by the visualization designer [Munzner, 2008]. This makes them uniquely challenging to visualize and physicalize, especially when multiple data variables are involved, as is common with current scientific data. Notwithstanding these factors, there is still a clear gap in the research for evaluation of spatial data physicalizations.

This chapter presents an evaluation of spatial data physicalizations compared with digital visualizations to address this gap. The work is critical to the visualization community because it contributes the first study in the literature which evaluates the effectiveness of physicalizations and digital visualizations in 2D and in virtual reality, for spatial data analysis tasks, as shown in Figure 6.1. Justified by prior studies, we hypothesize that physicalizations will – at a minimum – give performance measures *no worse* than the digital modalities, and that haptic perception and embodied interaction play a critical role in comprehending the data in the physical modality. To test all three modalities, participants complete three spatial data analysis tasks typical of those required in scientific sensemaking processes. We report the performance measures (time and errors) for each modality and task along with subjective measures (confidence and questionnaire feedback) provided by participants about each modality, and share the results which echo prior studies comparing modalities that physicalizations perform similarly to digital visualizations in most cases. Finally, we discuss the implications of data physicalization for spatial data analysis and scientific sensemaking processes.

6.2 Related Work

In this section, we introduce related work on evaluating the effectiveness of visualizations, and we build on the related work on data physicalization, tangible and embodied interactions, and extended reality visualizations from Chapter 2.

6.2.1 Evaluating Visualizations

For much of its history, visualization research was practiced as a craft rather than a robust scientific endeavor; a strong emphasis was placed on the novelty of new visualization techniques, while rigorous evaluation was less of a focus. Multiple calls to action were published in the early 2000s to turn the field's focus towards providing strong evaluations that validate the effectiveness of novel visualization techniques and systems [Kosara et al., 2003, Munzner et al., 2006]. In the two decades since, the field has seen a pronounced shift towards user study evaluations as a means to validate the effectiveness of visualizations: a Google Scholar search with the terms "*'user study' + visualization*" reveals only about 600 results in the year 2003 (the year Kosara et al.'s call to action was published), compared with approximately 5600 in 2023 (the time of writing of this dissertation).

Throughout the expansion of the user study as a means to validate new visualization techniques, multiple complimentary types of evaluation have emerged. Quantitative, objective values that are externally observed (e.g., time and errors) have been the primary measure for evaluating visualization techniques. Additionally, more qualitative, open-ended measures are available. Some qualitative measures impose a quantitative structure (e.g., a Likert scale), whereas others are assessed descriptively via coding or other means and these qualitative measures are subjective, because they are self-reported. Though the presentation of questions can be standardized across multiple experiments, participants' individual interpretations (and therefore responses) are not as verifiably repeatable across trials. Since there are so many evaluation methodologies now available to visualization researchers, it's imperative that the evaluation is designed to validate the intended contributions of a visualization design or work [Munzner, 2009]. The remainder of this section is devoted to exploring the related work in the evaluations this chapter builds on.

So far, evaluations in visualization research have primarily been focused on the performance, comprehension, and insights generated by new visualization techniques [Wang et al., 2019]. Several models have been proposed to assess the value of visualizations: these models take into account aspects of evaluation such as *effectiveness* and *efficiency* [van Wijk, 2005], and more specifically *time*, *insights*, *essence*, and *confidence* [Stasko, 2014]. Some studies focus explicitly on participants' time and errors while completing visualization tasks to evaluate different techniques (e.g., [Laidlaw et al., 2005, Zhang et al., 2016]). Other additional quantitative measures like number of insights generated are also common (e.g., [Saraiya et al., 2005, Yi et al., 2008]), as is measuring users' confidence in a visualization or visualization technique (e.g., [Wixon, 2011]). The work in this chapter builds on these styles of time, errors, and confidence evaluations to compare 2D, VR, and Physical data visualizations with spatial data analysis tasks.

While most visualization evaluations are currently based on quantitative performance measures, there has been an emerging body of research into evaluations beyond typical performance measures, stemming from such events as the over-a-decade-running BELIV workshop at the IEEE VIS conference [Bertini et al., 2008]. This work has prompted the development of additional conceptual models that broaden the *value* of visualization, and thus, the importance of new measures collected during evaluation that emphasize deeper, more meaningful results than pure quantitative performance data can provide. Saket et al. provide a review of visualization evaluations that focus on the user experience over strict performance measures, emphasizing such goals as memorability, engagement, and enjoyment [Saket et al., 2016]. Wang et al. propose an expanded model on the value of visualization that takes not only a visualization's intellectual contribution into account, but also the emotional, social, and physical engagement the visualization provides [Wang et al., 2019]. This model is especially relevant to evaluating data physicalizations; recall from Chapter 2 that Tangible and Embodied Interactions with data physicalizations can inherently provide these emotional, social, and physical values in visualization, by their nature of being physical objects. Though the evaluation presented in this chapter primarily focuses on the quantitative measures, the "intangible" benefits of tangible visualizations are quite clear. So, in the study presented in this chapter, we record additional qualitative measures that are better suited to capture the user experience and added value physicalizations contribute over raw *time* and *errors* performance measures.

Finally, it is essential the *tasks* used for evaluations are designed to accurately validate the visualization techniques used [Munzner, 2009]. Multiple taxonomies have deconstructed how data stakeholders use various types of visualizations and abstracted numerous tasks that are generalizable across application domains and datasets; here we describe two examples relevant to the work in the rest of this chapter. Brehmer and Munzner introduced a framework for classifying tasks based on their *what*, *why*, and *how* [Brehmer and Munzner, 2013]. This framework takes into account factors such as a user's motivation fo competing a task, the method they complete the task with, and the visualization components that enable task completion, and is validated through several examples and cross-examination with theoretical frameworks on information-gathering. More specific to spatial data visualization is the classification described by Laha et al., which generalizes tasks used for volumetric, spatial data [Laha et al., 2015]. This classification consists of six types of tasks commonly in diverse fields used to analyze spatial data, and is evaluated with an online survey of visualization practitioners who commonly use spatial data to generate new information and insights about their data. These two frameworks are used to describe the tasks in our study.

6.2.2 Comparing Physicalizations and Digital Visualizations

Taking inspiration from evaluations on digital visualizations, data physicalization researchers have also begun to develop robust evaluations comparing their effectiveness with other visualization techniques. The work in this chapter is inspired by two such studies.

Considered the first evaluation comparing data physicalizations to current 2D and 3D digital visualization modalities, Jansen et al. conducted a study which provided concrete evidence of the quantitative performance benefits that data physicalizations provide over digital data representations [Jansen et al., 2013]. Two studies were completed in the publication, one comparing physicalizations of 3D bar charts with digital 3D bar charts and 2D interactive linked-view bar charts, and one examining the role that touch plays in solving data analysis tasks using these bar charts. Statistical data from GapMinder[1] was used for both studies. In each study, participants completed three tasks: a *range* task to to find the minimum and maximum data values of a given group, a *compare* task to find the lowest value out of three given data points, and an *order* task to sort a group of given values from smallest value to largest value. The results found that participants performed significantly faster with 3D physical bar charts than they did with 3D digital bar charts, and they performed the fastest in the 2D interactive modality. Additionally, they concluded that the benefits of the physicalization are primarily from the embodied interactions and the act of physically *touching* and actively exploring the data object, rather than from just manipulating it.

Danyluk et al. introduced another study comparing participant performance on data physicalizations and in virtual reality [Danyluk et al., 2020]. Their study design was based on Jansen et al.'s original study with the same data and tasks, but the focus was on comparing multiple scales of virtual reality visualization with data physicalizations, with and without annotation capabilities. For all tasks and conditions, participants performed fastest in the physical modality, and additionally rated it the highest in terms of ease of use, speed, performance, and sharing with others. They concluded that while VR visualization tools are significant for their interactive annotation capabilities, current XR and digitally-driven haptics technology is a long way away from supporting the tactile and embodied interactions that physical objects inherently support.

Besides these two evaluations, few others exist that compare the performance and user experience of participants completing data analysis tasks, but we briefly describe other related studies that evaluate different metrics here. Prior to the invention of

[1] https://www.gapminder.org/

the term *data physicalization*, Tim Dwyer's Ph.D. thesis explored the visualization of abstract time-varying data in both physical and digital modalities, and showed that participants perform some tasks more efficiently in the physical modality [Dwyer, 2004]. Drogemueller et al. conducted a study comparing the performance and user experience of participants solving tasks on graph visualizations with four variations of visual and haptic feedback. They found that participants *perceived* themselves to be more accurate when they could see *and* touch the visualizations (physical), but did not detect any significant differences in performance [Drogemuller et al., 2021]. Ren and Hornecker conducted a study measuring participants' understanding and recollection of physicalizations and VR visualizations, and found that most participants tended to answer post-test questions faster when presented with the physicalization than with the VR visualization [Ren and Hornecker, 2021]. Novel display technologies like handheld spherical perspective-corrected displays (HPCDs) have also been compared with physical representations [Berard and Louis, 2017]. Although this study didn't use visualization-specific tasks, they found that participants performed faster with the digital HPCD than the physical object, though we note that participants were not able to *touch* the data, thereby likely losing some of the benefits of physical data objects [Dragicevic et al., 2021].

The study in this chapter takes inspiration from all these related works, and we primarily draw comparisons with on the studies by Jansen et al. and Danyluk et al. since they are the most relevant to the tasks used in this chapter.

6.3 Study Design

This section introduces a study design evaluating the effectiveness of participants performing analysis tasks on *spatial* data in three modalities: a *2D* visualization, a virtual reality (*VR*) visualization, and a data physicalization (*Physical*). The design, structure, and tasks of the study are all informed by two pilot studies with three participants each, as well as two prior studies from the literature [Jansen et al., 2013, Danyluk et al., 2020]. Though the prior studies inspire the study design, we do not attempt to replicate these earlier experiments due to the considerable differences necessitated by the switch from non-spatial to spatial data. Henceforth in this chapter, we will use "prior studies" to refer to these two specific papers.

6.3.1 Hypotheses

In this study, we test if participant performance measures in the Physical modality are – at a minimum – *no worse* than the other two modalities. Our reasoning for this is twofold. First, we know from the results of prior studies that participants perform certain tasks faster with physicalizations than digital visualizations [Jansen et al., 2013, Danyluk et al., 2020]. Second, we expect that performance measures are only the *beginning* of the advantages that encoding data in physical form presents; physical data objects can be more memorable [Cockburn and McKenzie, 2002], more relatable [Hornecker et al., 2023], and engage audiences in a way that digital data displays currently cannot [Wang et al., 2019]. So, *if participants' performance is no worse with spatial data physicalizations than with comparable digital visualizations, and physicalizations have considerable benefits beyond performance, then physicalizations are an option worth considering as a means for spatial data analysis.*

Thus, our main hypotheses are:

- **H1:** Participants' time measures are at least as fast in the Physical modality as in other modalities.
- **H2:** Participants' error measures are the same or lower in the Physical modality than in other modalities.
- **H3:** Participants rate their confidence at least as high in the Physical modality as in other modalities.

Though the primary focus of this study is on evaluating the comprehensibility of data physicalizations through performance measures (time and errors), the additional benefits of data physicalizations beyond performance are clear from the literature. Thus, we are motivated to identify a set of exploratory research questions to further investigate the outcomes of the study:

- **RQ1:** What role does touch play in the strategies used to complete spatial data analysis tasks in each modality?
- **RQ2:** What are the benefits and drawbacks of physicalization for completing spatial data analysis tasks?

Together with the main hypotheses, these research questions shape the design and analysis of this study comparing physicalizations, virtual reality and 2D desktop visualizations for spatial data analysis tasks. Given the additional advantages of physicalizations known to the literature (e.g., memorability, engagement, emotional connection, relatability, and embodiment), the support of these hypotheses and investigations into the research questions will motivate the adoption of data physicalizations in a variety of contexts.

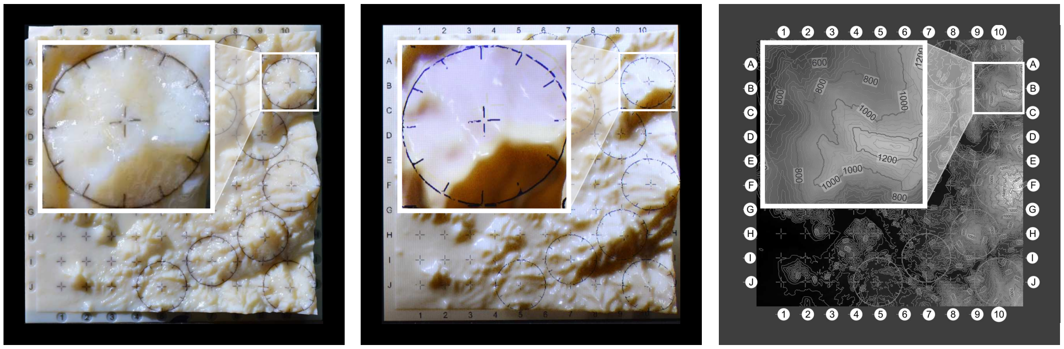

Figure 6.2: Visualization design for the Physical (left), VR (middle), and 2D (right) modalities. Insets show the details in each modality at the same location (circle B10). Note the color and contour lines in the 2D, and the light and shadows in the VR and Physical modalities.

6.3.2 Modalities and Visualization Design

Figure 6.1 shows the three modalities used in the study (Physical, VR, and 2D), and Figure 6.2 shows detailed views of the visualization design in each modality. When designing an evaluation for comparing multiple visualization modalities, one of the key challenges is to make the comparisons between modalities as equitable as possible (i.e., not giving an inherent advantage to any one modality) [Carpendale, 2008]. In the study design presented in this section, we proceed under the guiding principle that all modalities should be made as equal as possible. We evaluate all three modalities using digital elevation models (DEMs), each with a single scalar variable *elevation*.

Physical

Figure 6.2 (Left) shows the visualization design of the Physical modality. We start with the design of the physical visualization since it is the most constrained of the modalities (at least with current widely available digital manufacturing technology). In the Physical modality, we encode *elevation* using the physical object's height. In this way, it *resembles* a scaled model of the topography at the physical location where the data were recorded, but per our definition in Chapter 2, we note that it is still considered a *data physicalization* (not a model), contrary to some definitions of the term [Dragicevic et al., 2021].

Size plays a crucial role in the interpretability of data physicalizations [Jansen and Hornbaek, 2016, López García and Hornecker, 2021]. In our case, we desire to scale the physicalizations such that participants can pick them up; our pilot studies converged on a size of 15cm x 15cm to be large enough for participants to examine the intricate, high-resolution details of the spatial data but still small enough to hold, rotate, and touch. We also note that in prior studies participants were able to touch every bar in the physical 3D bar charts. Touching every part of the data becomes more complex with spatial data, especially since the data originate from real world elevations which tend to be bumpy and result in small crevasses where a finger cannot touch. As such, we add a constraint that the surface of every dataset must be at least 90% "touchable". The touchability constraint is verified during the data translation process; see Section 6.3.3 and Table 6.1 for more details.

The Physical modality also includes a physical legend to enable participants to look up elevations. The use of a legend is primarily motivated by the current constraints of digital fabrication: the physical data object is not of a sufficient scale to display labeled contour lines and have the text be readable, and using a continuous colormap was not feasible for the 3D printers available to us. The legend's numeric values are correct when the legend is resting anywhere on the flat margin surrounding the dataset. Every trial starts with the legend on the right margin of the physicalization, but the legend can be moved independently of the physical data object during trials.

The datasets for the Physical modality are 3D printed using a Stratasys Objet 500. This printer is limited in the colors it can reproduce, so a neutral white color (Rigur) is used for the base physicalization and the task overlay is black (Vero Black). The physical data representations each weigh about 0.5kg.

VR

Figure 6.2 (Middle) shows the visualization design of the VR modality. For the VR modality, we strive to create as close of a match as possible to the Physical modality for consistency. The interactions in the VR scene mimic those in real life, including rotation and translation, plus a bimanual scaling interaction to make the data larger or smaller. Using the Unity Engine's High-Definition Render Pipeline, we create physically-based materials and lights that match those from the Physical modality as closely as possible. For example, we match the positions, color temperature, and brightness of the lights, and match the index of refraction, color, and subsurface scattering diffusion profile of the physical material. However, despite these best efforts, the visualizations still do not look completely identical (see the comparison in Figure 6.2, note the difference in shadows). We attribute this to a combination of factors that have chronically plagued digital displays throughout their history, such as color reproduction issues, lack of true high dynamic range support, and light bouncing off objects in the real world that aren't fully modeled in the virtual scene.

Similar to the Physical modality, in VR we also encode *elevation* using the virtual object's height. The virtual object is initialized at a scale of 15cm x 15cm, but is resizable by the participant. The VR modality also includes a legend for looking up numeric elevations, and the same rules as the Physical modality apply regarding its positioning. However, since the virtual scene has no physics or gravity simulation, the virtual legend and virtual data object are positioned individually by the user without constraints.

The VR modality uses a flat-panel perspective-tracked stereoscopic display (known as a "fishtank" VR display [Ware et al., 1993]), coupled with bimanual spatial input devices. The dominant hand input device is a pen-shaped wand with an activation button; the non-dominant hand device is a simple textile loop with an activation button, a simplified single-stage version of the finger sleeve device described in [Zeleznik et al., 2002]. Virtual cursors are shown in VR for both hands, with the dominant hand cursor at the tip of the pen and the non-dominant hand cursor offset 5cm along the direction of the participant's finger so it appears approximately at their fingertip. Together, the head-tracked stereoscopic display and the input devices enable a participant to move their head around to see different viewpoints of the data object, and to "grab" the data object with their hands using the activation buttons.

The 3D display is a zSpace 300 with a screen measuring 61cm on the diagonal at 1920x1080 resolution, a 120Hz refresh rate (60Hz per eye), and polarized ocular separation. A fishtank-style VR display was used to give the VR modality the maximum possible resolution compared with the Physical modality. The spatial input devices consist of the zSpace built-in stylus and a Polhemus FASTRAK magnetic tracker, both of which provide position and rotation inputs at 60Hz. All photos of the VR modality were captured with a camera shooting through the zSpace tracked 3D glasses to obtain the correct viewpoint.

2D

Figure 6.2 (Right) shows the visualization design of the 2D modality. The 2D modality utilizes a typical 2D geographic information science (GIS) display of the data, and we follow current guidelines for visualizing such data (e.g., [Eynard and Jenny, 2016, Slocum et al., 2022]). Elevation is shown using a black-to-white colormap, where dark colors are low elevations and light colors are high elevations; labeled topographic contours are also shown at regular intervals along with five unlabeled sub-contours. The 2D visualization supports mouse-based pan-and-zoom interactions typical of map-based visualizations.

Consistency Across Modalities

We note that all three modalities are capable of much richer visualizations that encode multiple variables and include interactive filtering, querying, and annotation capabilities, which can make complex datasets more comprehensible [Shneiderman, 1996]. With the expanding capabilities of digital manufacturing technologies, it is now possible to encode increasingly complex data in physical form [Djavaherpour et al., 2021], and with recent work in augmented and active physicalizations (e.g., the approach to multi-touch querying in Chapter 5 and interactive shape-changing tables [Taher et al., 2016]), it is becoming progressively easier to introduce interactivity to data physicalizations. Prior studies comparing visualization modalities have included such functionality (e.g., interactive linked-view matrix plot and bar chart [Jansen and Dragicevic, 2013], annotation and highlighting [Danyluk et al., 2020]), but these advanced capabilities were only included for the digital modalities, *not* physical ones.

For this study, we return to the guiding principle that *to accurately compare modalities, visualization designs between modalities should be as equal as possible*. We reflect that the purpose of this study is *not* to compare interaction techniques or saliency of visual channels but to compare digital and physical modalities. Thus, we make the explicit design decision to explicitly use as few visual/haptic channels as possible to represent a single scalar *elevation* variable on the data, and for all modalities we only support basic interactions that manipulate the view of the visualization.

6.3.3 Data

The study uses high-resolution continuously sampled spatial data typical of those used in current data-intensive scientific sensemaking processes. Digital elevation models are real-world spatial datasets that scientists often use to develop a better understanding of the Earth's surface, for example, the watershed of a particular river [Millar et al., 2018], or the effects of land use on erosion [Tateosian et al., 2010]. This choice of data is in contrast to prior studies comparing visualization modality which use discrete, non-spatial data. Examples of the data represented in different modalities can be seen in Figures 6.1 and 6.2. While it is difficult to fully control for the "difficulty" of these data since they are from the real world, we chose datasets that had a variety of elevation ranges.

Table 6.1 shows the details of the four datasets used in the study, all of which originate from real elevation data near Sequoia and Kings Canyon National Parks, California, U.S.A. The data are downloaded from the United States Geological Survey open data portal as 1-degree-square GeoTIFF images, each at 3612×3612 resolution[2]. Then, the data are projected from their raw latitude/longitude into a Universal Transverse Mercator (UTM) projection to end up with a 2.5D heightmap with correct distances in meters. Lastly, the data are cropped into individual 24km x 24km squares, scaled down to 15cm

[2] https://apps.nationalmap.gov/downloader/, 1-degree 3DEP data

Dataset	UTM Coordinates	Elevation Range	% Valid Points
0	-13247189.5, 4300449.7	1778m	96.9%
1	-13172667.3, 4300449.7	1189m	96.5%
2	-13209928.4, 4300449.7	1860m	91.6%
3	-13247189.5, 4369884.8	2694m	93.4%

Table 6.1: Four datasets are used for the study; dataset 0 is used for training only. The upper-left latitude/longitude coordinates in a Universal Transverse Mercator projection, the total elevation range, (maximum elevation - minimum elevation), and the percentage of valid points in the dataset that are touchable are provided for each dataset.

x 15cm squares, and translated into 3D using Blender 3D modeling software and the Blender GIS[3] plugin. For the data shown in both 3D modalities, we apply a vertical exaggeration of 5x to the height data, as is common practice in scientific visualization (e.g., [Mitasova et al., 2012]); this leads to datasets ranging from 3cm - 7cm tall. During the translation step, we also verify the "touchability" of each dataset by sweeping an approximated fingertip (a 10mm radius paraboloid, $f(r) = \frac{1}{10}r^2$) over the surface and checking to see if the "finger" intersects any other point on the surface. Any dataset that had more than 10% of points where the finger intersected another point on the surface was discarded.

Every dataset in every modality also includes an overlay which consists of a 10×10 grid of crosshairs marking points across the data; rows are labeled by letters A-J and columns are labeled by numbers 1-10. Additionally, 10 circular "clocks" divided into 12 sections are included on the task overlay to provide enclosed areas and circular angle references; 12 o'clock noon is North on every clock (pointing toward the top of the dataset). The overlay geometry is created with a Blender Python script during the data translation process. The components of this overlay together enable participants to complete tasks relating to specific points and areas in the spatial data without repeating the same task with any given point.

6.3.4 Tasks

Figure 6.3 shows the tasks used in this study. These three tasks are representative of real-world questions a scientist might ask about spatial data, for example understanding the gradient of steepest descent in a particular area (*Advect*), comparing elevations at multiple data points (*Compare*), and finding the minimum and maximum elevation in a particular area (*Range*). For all tasks participants are given a *target* to solve the task on, identified by its row/column ID as shown on the task overlay in Figure 6.3. The

[3] https://github.com/domlysz/BlenderGIS/tree/master

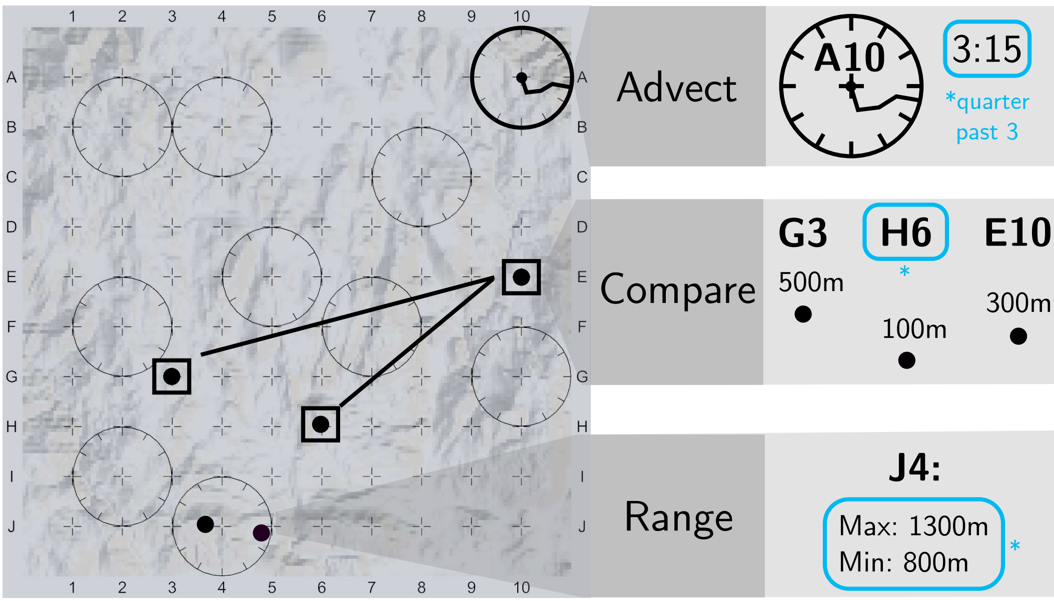

Figure 6.3: A grid of crosshairs and circles is overlaid on the elevation data to describe targets for each task. Tasks include an Advect task to trace a path downhill from the center of a target circle to the edge of the circle, a Compare task to determine the point with the lowest elevation out of three target points, and a Range task to determine the minimum and maximum elevation inside a target circle. Example ground-truth answers notated with an asterisk *.

three tasks are not trivial to complete, yet they are well defined in terms of task taxonomies [Brehmer and Munzner, 2013, Laha et al., 2015], and ground-truth answers are programmatically calculable; this makes them ideally suited to this study. Additionally, the Compare and Range tasks are adapted from prior studies, which will enable us to cross-examine the similarities and differences in the effectiveness of data physicalizations with and without spatial data.

The exact wording for the tasks given to participants is:

- **Advect**: Trace the path downhill from a given point until the path intersects the surrounding circle. What heading (o'clock) does the path intersect the circle?
- **Compare**: Locate the three given points and determine which point has the lowest elevation.
- **Range**: Indicate the range of elevations inside a circle surrounding a given point.

To complete the Advect task, participants are asked to trace the path of steepest descent from the center of a target circle to the edge, and report the o'clock heading where the downhill path intersects. This task is typical of data-intensive scientific sense-making processes, and it can be described by Laha et al.'s classification as an *absolute, spatial understanding* task. Participants are required to first search the dataset for the target circle, then to develop a deep understanding of not only the 3D elevation data within an area, but also its gradient of steepest descent. During the task introduction, participants are instructed to imagine a tiny ball rolling downhill starting at the center of the circle as an analogy to help with task comprehension. Participants are directed to report their answers to the nearest quarter-hour (e.g., quarter past 3), which equates to an angular resolution of 7.5°. This task is inspired by prior studies on fluid flow visualization (e.g., [Laidlaw et al., 2005, Forsberg et al., 2009]).

To complete the Compare task, participants are asked to locate three target points by their IDs and report which one has the lowest elevation. This task is typical of those found in many information-seeking contexts in visualization; the motivation for the task can be described using Brehmer and Munzner's framework as *locate* and *compare*, and it is an example of a *relative spatial understanding* task in Laha et al.'s classification. Participants are required to locate each of the three points on the visualization, then to make relative judgments of their elevations, and report their answer as the lowest point's ID.

To complete the Range task, participants were are to locate a target circle and report its minimum and maximum elevation. This is another typical task seen across many contexts in the literature. By Brehmer and Munzner's framework, it can be described as a *locate*, *lookup*, and *identify* task, and in Laha et al.'s classification it can be considered

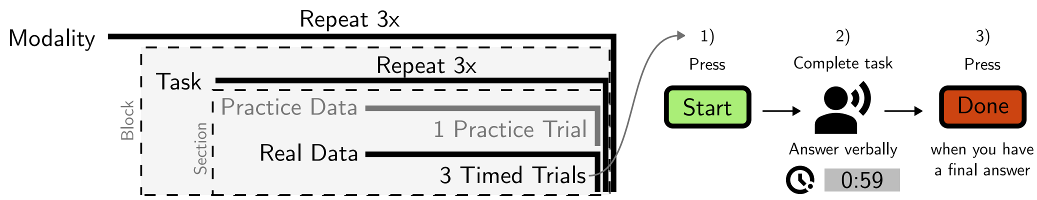

Figure 6.4: Procedure for the study. Participants complete 9 practice trials and 27 timed trials. Each block contains 3 modalities, and each section contains 3 tasks, each with 4 trials (1 practice, 3 timed). In each trial, participants begin the trial by pressing the "Start" button on the tablet, complete the task and state the answer verbally, then press the "Done" button when they have a final answer.

an *absolute quantitative estimation* task. Participants are required to first locate the target circle by its ID, then estimate and report the numeric minimum and maximum elevation of the area within the circle.

Ground-truth task answers like those shown in Figure 6.3 are calculated directly from the 3D DEM mesh data using a Blender Python script during the data translation process. For the Advect task, the ground-truth path of steepest descent is calculated by starting with the vertex at the center of the target circle and iteratively finding the steepest gradient for all connected vertices until encountering the edge of the circle; circles with a local minimum are excluded. For the Compare task, the ground-truth answer is calculated as the minimum elevation of the three given points; each point uses the elevation at the closest vertex in the DEM. For the Range task, the ground-truth answer is calculated by finding the two vertices in each circle in the DEM with the minimum and maximum elevation.

6.3.5 Experimental Procedure

Figure 6.4 shows the procedure that participants follow through the study. We use a full-factorial within-subjects design where every participant is exposed to 3 *Modalities* $\times$ 3 *Tasks*, each with 1 practice trial and 3 timed trials, so participants all complete a total of 9 practice trials and 27 timed trials in the study. The presentation order is counterbalanced with two independent balanced Latin squares, one for task and one for modality. Participants are given the same dataset for all timed trials, which is one of the datasets 1-3 shown in Table 6.1, and a separate dataset is given for practice trials (dataset 0). The task overlay is designed to avoid repeating the same target more than once for the same task in the duration of the experiment.

After informed consent, a short introduction to the study is given, including an overview of the modalities used in the study, the task overlay notation with a quiz to check participants' understanding of the notation, and instructions on how to proceed through the study using the tablet interface. Additionally, we instruct participants to do their best to respond to the tasks within one minute or so; the exact text used for this section is:

> *Try to answer the prompt in <1 minute. Audio and visual prompts will be provided at 0:30, and 1:00 onwards. If you find yourself beyond the 1:00 mark, try to finish up. Give it your best answer; no need to be perfect!*

We added this time constraint because in multiple pilot studies, some participants tended to take more than 5 minutes on certain tasks trying to make their answer perfect, and this was problematic for minimizing fatigue effects. We found that one minute was a sufficient amount of time for participants to converge on a reasonable level of accuracy. With this constraint in place, participants took about 75 minutes on average to complete the study.

Beginning each block, participants are introduced to each modality. Specifically, in the VR modality we instruct participants how to use the bimanual controls to move, rotate, and scale the virtual data object, and in the 2D modality we instruct participants how to use the mouse buttons to pan and zoom the topographic map. Based on our pilot studies, the Physical modality needs only a brief introduction stating that participants are welcome to pick up, touch, hold, and manipulate the physical data object and legend. At the end of each block, participants are given a questionnaire, and are instructed to take a short break in between blocks.

At the beginning of each section, participants are given instructions on how to complete the task in the current modality. During each section, participants complete 1 practice trial with a practice dataset, then 3 timed trials with a "real" dataset. Figure 6.4 (Right) shows the process participants follow in each trial: 1. Press "Start" button, 2. Report answer verbally, and 3. Press "Stop" button when you have a final answer. After each section is complete, participants are asked to report their confidence for that section (modality and task pair).

Measures

For each trial, we record two objective, quantitative measures: the *time* participants took and the *errors* participants committed. Time is measured as the time difference between the participant pressing the "Start" button to begin the task and the "Done" button to signify they had a final answer. Error calculations differ per-task, and all represent normalized values 0 (correct) to 1 (incorrect). For the Advect task, error is calculated as the percentage of a full circle away from the correct answer (e.g., 180°= 0.5). For the Compare task, a correct answer is reported as 0, and an incorrect answer as 1. For the Range task, error is calculated in Equation 6.1 as the absolute difference between minimum and maximum elevations the user answered and the ground truth range divided by the full data range of the dataset.

$$\frac{(|\mathrm{min}_{user} - \mathrm{min}_{correct}| + |\mathrm{max}_{user} - \mathrm{max}_{correct}|)/2}{\mathrm{max}_{data} - \mathrm{min}_{data}} \tag{6.1}$$

To obtain further insights into participants' user experience and task-solving strategies, we also record multiple subjective measures. First, participants report their *confidence* at the end of each section (modality and task pairing). Specifically, a prompt on the tablet asks participants *"On a scale from 1 to 5, where 1 is least confident and 5 is most confident, how confident were you in your answers for the section you just completed?"*. Then, at the end of each block (modality), participants fill out a questionnaire with free-response questions:

- Were there any positive factors that contributed to your experience with the user interface? If so, please describe them.
- Were there any negative factors that took away from your experience with the user interface? If so, please describe them.
- Were there any strategies that helped you complete the task? If so, please describe them.

All sessions are recorded with a video camera to analyze how participants used each modality by touch and manipulation.

Setup

Figure 6.5 shows the physical setup for the study, including a 3D display, a data physicalization and its corresponding legend set atop a square base, a mouse, a tablet, polarized perspective-tracked 3D glasses, and bimanual spatial input devices. The lighting in the room is controlled for the experiment: two diffuse lamps were mounted above and to the sides of the work area, each 1200 lumens with a color temperature of 5000K; there is an additional 180 lumen spot lamp to the left side pointing at the work area which had a color temperature of 4800K. All prompts are displayed on the tablet; participants also start and end each trial using buttons on this tablet. Except for breaks in between each block, participants remain seated at the same position in front of the work area for the duration of the study.

For the Physical modality, the physical data object sits on its base in front of the user, the computer's input devices are removed from the desk, and the screen on the 3D display is blanked. For the VR modality, the physical dataset is removed from the

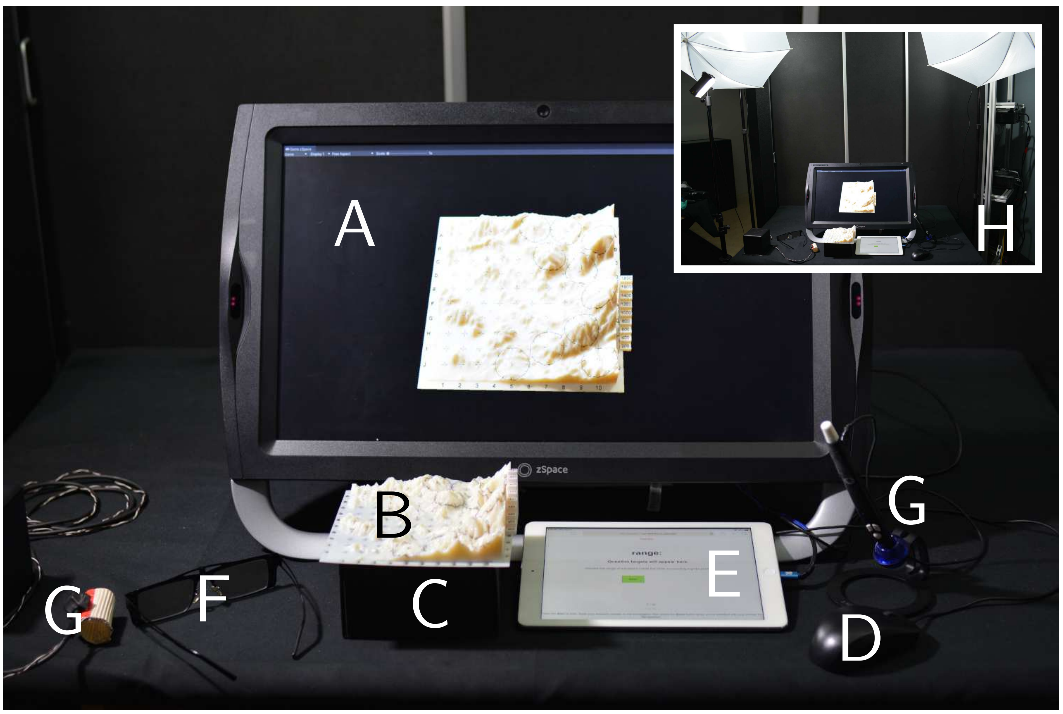

Figure 6.5: Experimental setup for the study: A. 3D display; B. Data physicalization and legend; C. Physicalization base; D. Mouse; E: Tablet; F. 3D Glasses; G. Spatial input devices; H (inset): Controlled lighting sources. Not pictured: video camera.

desk; 3D glasses are used to provide a perspective-tracked stereoscopic view of the virtual visualization, and the spatial input devices are used to move, rotate, and scale the virtual object. For the 2D modality, the same display is used with 3D mode turned off; only the mouse remains on the desk enabling participants to pan and zoom the 2D visualization.

Participants

The study included 24 participants recruited from a university student population and a local musical ensemble via electronic newsletter and posted fliers. Multiples of 6 participants were necessary to ensure the validity of the balanced Latin squares for modality and task, and we used G*Power [Faul et al., 2009] to determine the sample size *a priori* with a *medium* effect size of 0.5. Thus, the study had *24 Participants* $\times$ *3 Modalities* $\times$ *3 Tasks* $\times$ *3 Trials* $=$ *648 Total Trials*. There were 10 Female and 14 Male participants, ages ranging from 19 to 52. Three participants were left-handed, and all participants reported having normal or corrected-to-normal vision. 5 participants had never used VR before this study, and 6 had never used topographic maps (*VRUsage* $=$ *have not used*, and *TopoUsage* $=$ *have not used*, respectively). At the end of the study, participants were compensated with a $15 USD electronic gift card.

6.4 Results

This section reports results for the evaluation comparing 2D, VR, and Physical visualizations of spatial data. We give an overview of the results based on our hypotheses and research questions identified in Section 6.3.1; we include results on the comprehensibility of the data via objective quantitative performance measures (time and error), as well as further subjective measures including quantitative confidence and qualitative feedback from participants.

6.4.1 Quantitative Measures

Figure 6.6 summarizes the quantitative results on the *time*, *errors*, and *confidence* participants had in each *Task* and *Modality*. Quantitative results are reported for each measure, and results are separated by task, as performance and confidence between tasks are quite different and comparing them is generally not useful.

To test the main hypotheses presented in Section 6.3.1, we used a two one-sided test procedure (TOST) to test equivalence/non-inferiority of the Physical modality to the other modalities [Walker and Nowacki, 2011]. This approach is better suited than traditional superiority testing (i.e., repeated-measures ANOVA and pairwise t-tests) to

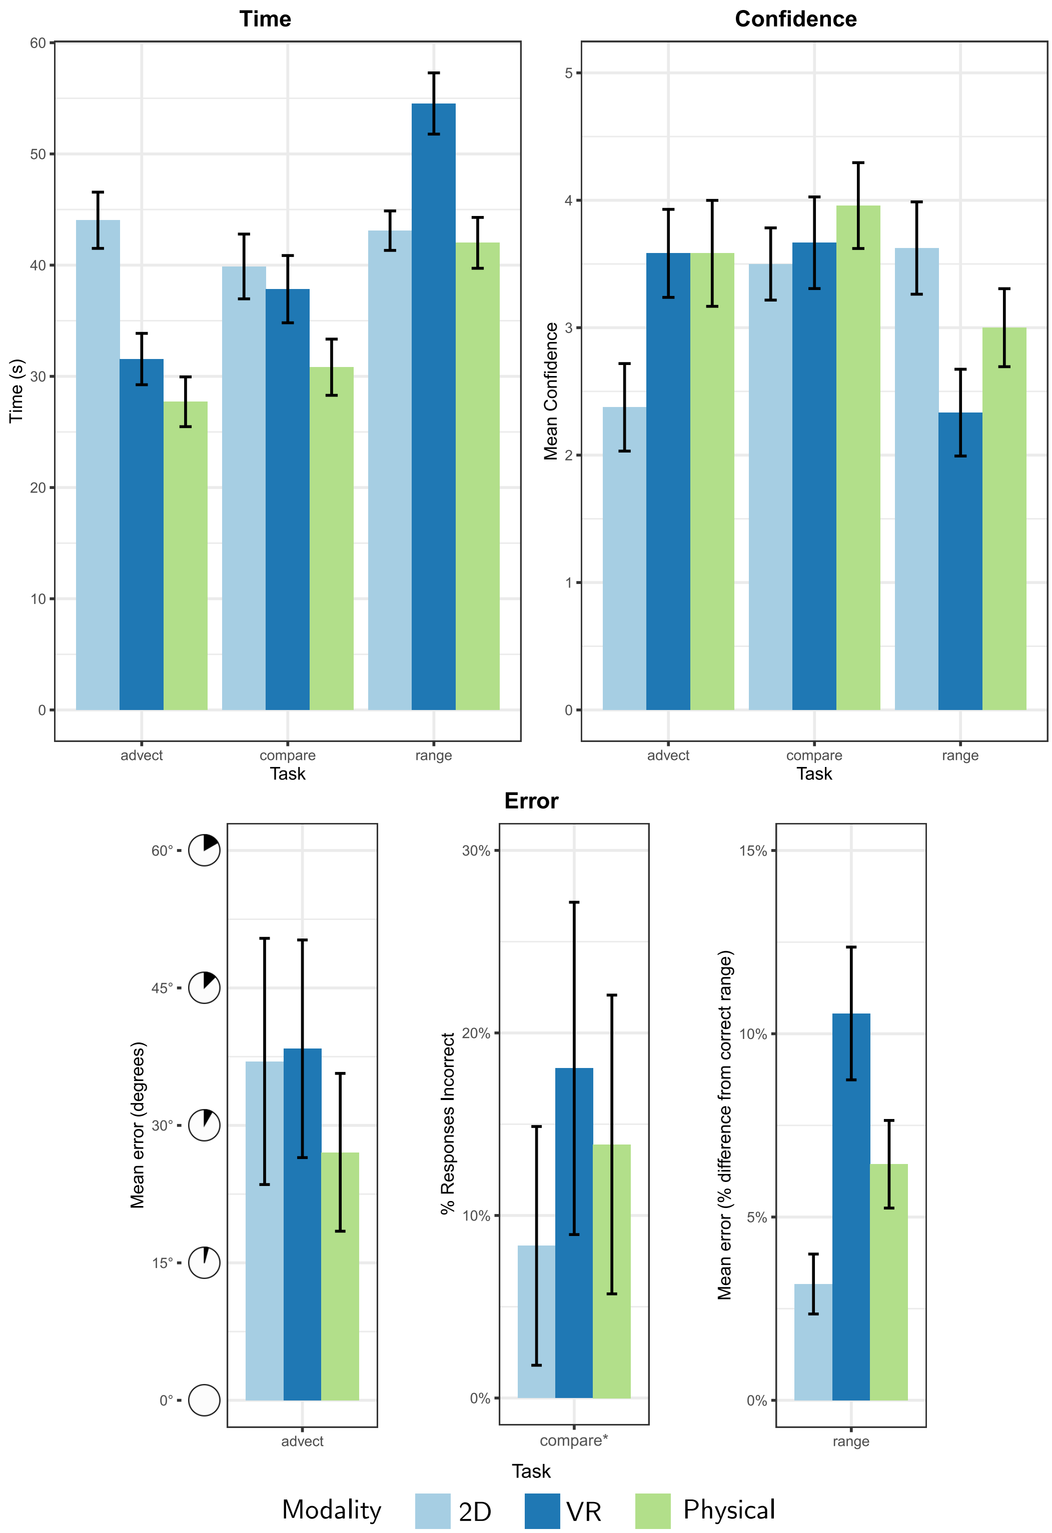

Figure 6.6: Time, errors, and confidence results for the study. Error bars represent 95% correlation-adjusted confidence interval [Cousineau, 2019]. **Compare Error uses standard 95% CIs because the assumptions of sphericity and compound symmetry.*

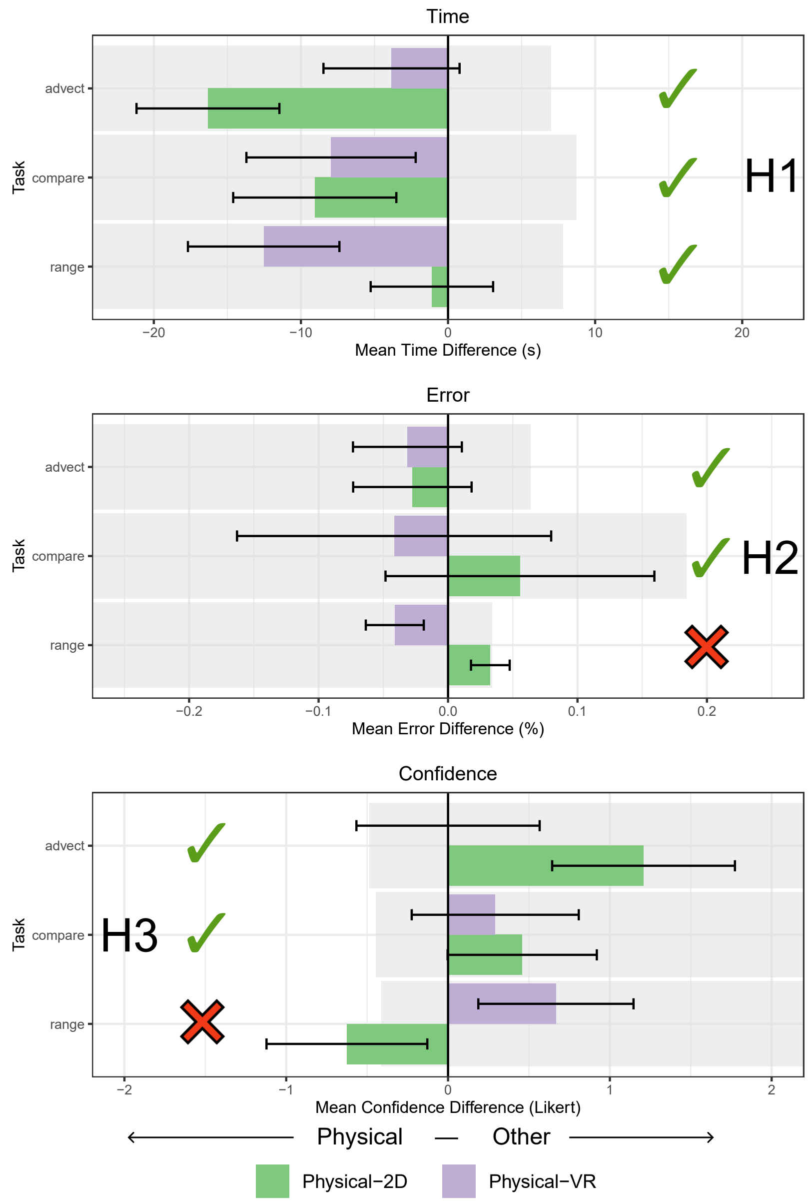

Figure 6.7: Pairwise *time*, *error*, and *confidence* differences between modality (*Physical* − *Other*, 0 indicates no difference between modalities). The shaded regions indicate the hypotheses from Section 6.3.1 (with the equivalence margin $\delta = 0.5$, i.e. $0.5\sigma_{pooled}$ from the differences), CIs completely overlapping the shaded region indicates support for the hypothesis. Confidence intervals indicate 95% CI calculated from mean differences.

test the hypotheses that Physical gives performance measures *no worse* than either other modality. When using equivalence and non-inferiority tests, a critical decision is the equivalence margin, also known as the smallest effect size of interest. The equivalence margin for this study was determined using the standard mean difference effect size from pilot study results, calculated using Cohen's d [Cohen, 2013] and bias-corrected by Hedges' g [Hedges and Olkin, 2014]. We used symmetric upper and lower equivalence margin bounds for the TOSTs as $\delta_{lower} = -0.5$ and $\delta_{upper} = 0.5$ (i.e., 0.5σ); this also corresponds with the *medium* effect size used in our *a priori* power analysis to determine the sample size. All TOST calculations were performed using the R package `TOSTER`[4] and the approach described by Lakens et al [Lakens et al., 2018]. Figure 6.7 shows pairwise differences between each modality for each measure with the hypotheses from Section 6.3.1 shaded at 0.5σ to correspond with the equivalence margin used in the TOSTs.

Except where noted, all confidence intervals refer to the 95% correlation-adjusted CI appropriate for comparing groups in within-subjects designs, where the assumption of sphericity and compound symmetry holds on the data [Cousineau, 2019]; CIs are notated by square brackets in the text. For the *time* and *error* measures except Compare error, the assumptions of log-normality and equality of variances held, so the log-transformed measures were used for analysis.

Time

Figure 6.6 (top left) shows the absolute time results for each task, and Figure 6.7 (top) shows the relative differences in time between the 2D and VR modalities compared with the Physical modality, for each task; the shaded boxes for each task show the non-inferiority equivalence margin of $\delta = 0.5\sigma$.

For the Advect task, the Physical modality was fastest on average (mean 27.7s, 95% correlation-adjusted CI [25.5s, 0.0s]), followed by VR (31.6s [29.2s, 33.9s]) and the 2D modality was slowest (44.0s [41.5s, 46.5s]). The upper-bound non-inferiority TOST indicated that Physical was not significantly slower than VR ($t(71) = -13.37, p < 0.001$) and Physical was not significantly slower than 2D ($t(71) = -16.53, p < 0.001$). This result supports hypothesis **H1** for the Advect task.

For the Compare task, Physical was fastest on average (30.8s [28.3s, 33.4s]), followed by VR (38.7s [35.7s, 41.9s]), and 2D was the slowest (39.9s [37.0s, 42.8s]). The upper-bound non-inferiority TOST indicated that Physical was not significantly slower than VR ($t(71) = -12.84, p < 0.001$) and Physical was not significantly slower than 2D ($t(71) = -14.51, p < 0.001$). This result supports hypothesis **H1** for the Compare task.

[4] `https://aaroncaldwell.us/TOSTERpkg/`

For the Range task, Physical had time (42.0s [39.7s, 44.3s]), 2D was (43.1s [41.3s, 44.9s]), and VR was (54.5s [51.8s, 57.3s]). The upper-bound non-inferiority TOST indicated that Physical was not significantly slower than VR ($t(71) = -18.71, p < 0.001$) and Physical was not significantly slower than 2D ($t(71) = -12.49, p < 0.001$). This result supports hypothesis **H1** for the Range task.

Besides the results from the main hypotheses, we report additional important results for the *time* measure. First, we consider the effects of VR usage and topographic map usage prior to this study. A repeated measures ANOVA on *time* revealed an effect of *VRUsage*, but only for the VR modality ($F(1, 214) = 7.40$, $p < 0.01$). In VR, participants who had previously used VR were significantly faster (40.0s [37.1s, 42.9s]) versus those who had not (47.9s [52.2s, 53.6s]). A repeated measures ANOVA on *time* also revealed an effect of *TopoUsage*, but only on the 3D modalities. Participants who had previously used topographic maps were significantly faster in both 3D modalities (Physical: 31.6s [29.3s, 33.9s]; VR: 38.7s [35.8s, 41.6s]) versus those who had not (Physical: 39.2s [34.5s, 43.8s]; VR: 50.4s [45.1s, 55.7s]), but 2D did not see such differences (have used: 42.0s [39.7s, 44.2s]; have not used: 43.5s [39.0s, 48.0s]).

We also evaluate the *time* measure with respect to each *Dataset* shown in Table 6.1. A repeated measures ANOVA on *time* detected an effect of *Dataset* ($F(2, 645) = 37.949$, $p < 0.001$). Post-hoc pairwise t-tests revealed significant differences between all pairs of datasets. Participants performed the fastest with Dataset 3 (33.3s [31.1s, 36.5s]), followed by Dataset 2 (37.7s [36.0s, 39.6s]), with Dataset 1 being the slowest (46.4s [43.9s, 49.0s]).

Errors

Figure 6.6 (bottom) shows the absolute error results for each task, and Figure 6.7 (middle) shows the relative differences in error between the 2D and VR modalities compared with the Physical modality, for each task; the shaded boxes for each task show the non-inferiority equivalence margin of $\delta = 0.5\sigma$.

For the Advect task, Physical had an average angular error of (27° [19°, 37°]), compared with 2D (37° [24°, 50°]) and VR (38° [27°, 50°]). The upper-bound non-inferiority TOST indicated that Physical did not have significantly higher error than VR ($t(71) = -3.40, p < 0.001$) and Physical did not have significantly higher error than 2D ($t(71) = -2.09, p < 0.05$). This result supports hypothesis **H2** for the Advect task.

For the Compare task, 2D had an average incorrect response rate of (8.3%, [1.8%, 14.9%]), Physical was (13.8% [5.7%, 22.0%]), and VR was (18.1% [8.9%, 27.2%]); the reported CIs are standard 95% CIs because the assumptions of sphericity and compound symmetry do not hold for Compare error. Additionally, since the Compare error was not normally or log-normally distributed, we use a Wilcoxon signed rank TOST instead of a standard t-test TOST. The upper-bound non-inferiority TOST indicated that Physical did not have significantly higher error than VR ($Z = 504, p < 0.001$) and Physical did not have significantly higher error than 2D ($Z = 615, p < 0.001$). This result supports hypothesis **H2** for the Compare task.

For the Range task, 2D had the lowest error on average (3.6% [2.4%, 4.0%]), followed by Physical (6.4% [5.2%, 7.6%]), followed by VR (10.6% [8.7%, 12.4%]). The upper-bound non-inferiority TOST indicated that Physical did not have significantly higher error than VR ($t(71) = -7.28, p < 0.001$). However, lower-bound non-inferiority test indicated that 2D had significantly lower errors than Physical ($t(71) = 8.88, p < 0.001$). This result does not hypothesis **H2** for the Range task.

A repeated measures ANOVA on *error* did not find an effect of *VRUsage* or *TopoUsage*, nor did it find an effect of *Dataset*.

Confidence

Figure 6.6 (top right) shows the absolute confidence results for each task, and Figure 6.7 (bottom) shows the relative differences in confidence between the 2D and VR modalities compared with the Physical modality, for each task; the shaded boxes for each task show the non-inferiority equivalence margin of $\delta = 0.5\sigma$.

For the confidence analysis, we take an "intervalist" interpretation to analyzing the 5-point Likert scale data, thus we treat the measure as continuous [South et al., 2022]. This enables us to use the arithmetic mean to perform analyses, which is considered a truer descriptor than the median when summarizing Likert scale data [Lewis, 1993].

For the Advect task, participants were most confident in 3D, with Physical and VR being equal (Physical: 3.58 [3.17, 4.00]; VR: 3.58 [3.24, 3.93]), and they were least confident in 2D (2.38 [2.03, 2.72]). The lower-bound non-inferiority TOST indicated that Physical did not have significantly lower confidence than VR ($t(23) = 2.22, p < 0.05$), and Physical did not have significantly lower confidence than 2D ($t(23) = 5.92, p < 0.001$). This result supports hypothesis **H3** for the Advect task.

For the Compare task, participants ranked their confidence in Physical (3.96 [3.62, 4.30]), VR was (3.67 [3.31, 4.03]), and 2D was (3.50 [3.22, 3.78]). The lower-bound non-inferiority TOST indicated that Physical did not have significantly lower confidence than VR ($t(23) = 4.52, p < 0.001$), and Physical did not have significantly lower confidence than 2D ($t(23) = 4.42, p < 0.001$). This result supports hypothesis **H3** for the Compare task.

For the Range task, participants were most confident in the 2D modality (3.63 [3.26, 3.99]), followed by Physical (3.00 [2.69, 3.31]), followed by VR (2.33 [1.99, 2.67]). The

Task	Manipulate		Touch	
	Strategy	**Frequency**	**Strategy**	**Frequency**
Advect	None	71%	None	31%
Advect	Rotate	4%	One Finger	57%
Advect	Pick Up	25%	Multiple Fingers	12%
Compare	None	58%	None	7%
Compare	Rotate	1%	One Finger	19%
Compare	Pick Up	40%	Multiple Fingers	74%
Range	None	39%	None	40%
Range	Rotate	19%	One Finger	36%
Range	Pick Up	42%	Multiple Fingers	24%
Mean	*None*	*56%*	*None*	*26%*
Mean	*Rotate*	*8%*	*One Finger*	*37%*
Mean	*Pick Up*	*36%*	*Multiple Fingers*	*37%*

Table 6.2: Touch and manipulation in the Physical modality

lower-bound non-inferiority TOST indicated that Physical did not have significantly lower confidence than VR ($t(23) = 4.74, p < 0.001$). However, the upper-bound TOST indicated that 2D had significantly higher confidence than Physical ($t(23) = -5.22, p < 0.001$). This result does not support hypothesis **H3** for the Range task.

Touch and Manipulation

Table 6.2 shows the frequency that participants touched and manipulated the physicalization in the Physical modality trials, captured by inspecting the video recordings for each trial. Starting with manipulation, we note that during a majority of trials, participants did not opt to pick up or rotate the physicalization on its base, though it varies by task. For the Range task, participants picked up the physicalization in just under half the trials, and rotated it in just under a quarter. For the Compare task, the behavior was split between no manipulation and picking up the physicalization. For the Advect task, just under a quarter of participants picked up the physicalization while the rest left it on its base.

With touch, we observe that during three quarters of the trials, participants chose to touch the physicalization with at least one finger (Compare = 93%, Advect = 69%, Range = 63%; *mean = 75%*). Breaking it down further, most participants used one finger for the Advect task and multiple fingers for the Compare task, while the touch strategies for the Range task were more mixed. Every participant touched the physicalization during at least one trial. Besides the touch usage in the Physical modality, we note that some participants also used touch as a strategy in 2D; during 7% of trials

2D		VR		Physical	
color	10	manipulate	14	physicalization	14
contour lines	8	virtual visualization	13	manipulate	9
manipulate	8	input devices	2	touch	9
2D visualization	6	legend	2	legend	3
labels	4	lighting	2	confirmation	1
cursor	3	bounds	1	lighting	1
background	1	clipping	1	points	1
		cursor	1		
sum	*40*	*sum*	*36*	*sum*	*38*

Table 6.3: Positive factors identified by participants.

2D		VR		Physical	
slow	13	legend	15	legend	10
labels	8	manipulate	13	labels	2
manipulate	8	tracking	3	lack of resources	2
contour lines	6	input devices	2	solid	2
crosshairs	2	virtual visualization	2	color	1
color	1	clipping	1	size	1
sum	*38*	*sum*	*36*	*sum*	*18*

Table 6.4: Negative factors identified by participants.

participants used multiple fingers and 10% of trials they used a single finger.

6.4.2 Qualitative Feedback

To better understand the participants' user experience during the study, we analyzed the qualitative subjective answers to free-response questions. The open-form text from participants' feedback responses was coded using the questions we asked as a codebook, including positive and negative factors, and strategies participants used (see Section 6.3.5). Two coders independently distilled the responses from each question; they were not given any information about responses beyond the question text to minimize bias in the coding. The coded responses were then normalized for meaning; for example, "lines" would be normalized to "contour lines", and "feel" would be normalized to "touch".

2D		VR		Physical	
color then contour lines	6	manipulate	13	touch	11
contour lines	3	change view	7	change view	9
cursor	3	align legend	4	manipulate	7
trace path	3	clipping	1	paired comparison	2
color	2	cursor	1	trace path	2
count contour lines	2	lighting	1	align legend	1
manipulate	2	picture ball rolling	1	connect points	1
map context	2	recheck answers	1	legend	1
paired comparison	2	trace path	1	lighting	1
touch	2			picture ball rolling	1
labels	1			picture water flowing	1
picture ball rolling	1				
shading	1				

Table 6.5: Strategies participants used during the study.

Positive and Negative Factors

Table 6.3 shows the coded positive factors reported by participants in their qualitative feedback. For all three modalities, participants frequently mentioned both the fact that they could *manipulate* every visualization, and the most salient components of each modality. In both 3D modalities, the visualization (*virtual* or *physical*) was mentioned by a majority of participants, as were the *color* and *contour lines* in 2D. Nine participants also mentioned *touch* as a positive factor while using the Physical modality. Overall, there were an approximately equal amount of positive factors for each modality.

Table 6.4 shows the coded negative factors reported by participants in their qualitative feedback. For the 2D modality, a majority of participants mentioned that the using the mouse to *manipulate* the viewpoint was *slow*, and that *labels* were sometimes hard to read. In VR, a majority of participants found it difficult to *manipulate* the virtual data object with the spatial *input devices*, and three participants mentioned some *tracking* issues. For the Physical modality, two participants mentioned the lack of *labels* and the fact that the physical data object was *solid* as negative factors. In both 3D modalities, the most common negative factor was the *legend*, with 15 participants mentioning it for VR and 10 for Physical. Finally, participants identified fewer than half the negative factors in Physical that they did in either digital modality

Strategies

Table 6.5 shows the coded strategies reported by participants in their qualitative feedback. Participants in all three modalities expressed a variety of strategies. Strategies in

2D were the most varied, with some participants first looking at *color*, then at *contour lines*, while others only looked at the *contour lines*, or used a *cursor* or *touch* to mark points. In both 3D modalities, *manipulate* was a predominant strategy, as was *changing view* to look around the visualization. For the Physical modality, eleven participants mentioned *touch* as an important strategy to completing the tasks, and in VR, four explicitly mentioned they needed to *align legend*.

6.5 Discussion

In this section, we discuss the results of the study in the context of both the quantitative and qualitative feedback, and we compare the results of this study with prior studies from the literature.

6.5.1 Interpretation of Results

Overall, the results for the time measures support hypothesis **H1**: in all tasks, the Physical modality was at least as fast the as other modalities. Support for hypothesis **H2** is divided by task: the Physical modality shows errors equal or less than other modalities in the Advect and Compare tasks, and shows higher errors than 2D in the Range task. Similarly for hypothesis **H3**, participants report similar or higher confidence ratings for Physical versus other modalities in the Advect and Compare tasks, and report lower confidence in Physical than 2D for the Range task. The rest of this section discusses the results and implications for each hypothesis.

Time and Errors

For the Advect task, the Physical modality was the fastest, without incurring higher errors. Indeed, there was not an observable speed/accuracy trade-off between modalities for this task; even though participants took longer in 2D and VR modalities, their mean error was actually *higher* than in the Physical modality. Returning to the motivation for using physical data objects, we attribute the success of the Physical modality for the Advect task to three main factors: the embodied role of touch on the physical data object, the haptic perception of the data, and the real-world depth perception that a physical object provides. We suspect these advantages are critical to the higher performance of the Physical modality, particularly for this task which requires participants to develop a deep understanding of not only the surface of the elevation data but also its gradient.

For the Compare task, the Physical modality was faster than the other two modalities. We also observe that participants committed the least errors in 2D; there is a noticeable speed/accuracy trade-off for this task. From the qualitative feedback and video recordings, we observed that people used very different strategies for the 2D and the Physical modalities; in 2D, nine participants reported that they counted individual contour lines, whereas in the Physical multiple participants reported that they used their finger to mark points, and make quick pairwise comparisons between points' elevations. These varied strategies have also been reported in the literature; physicalizations, particularly those that are hand-scale and able to be picked up, tend to lend themselves naturally to making quick judgments [López García and Hornecker, 2021]. So, we suspect this difference in strategies afforded by each modality was at least partially responsible for the speed/accuracy trade-off seen in the Compare task.

For the Range task, the Physical and 2D modalities had about the same time on average, but 2D had fewer errors; VR had both higher time and more errors for this task. In both the Physical and the VR modalities, participants were required to look up the numeric elevation on the separate legend, which participants found to be both cognitively demanding and inaccurate. In fact, the *legend* was the most common negative factor for both 3D modalities. This was not an issue in 2D, where the numeric elevations were provided as labeled contour lines. Like the Compare task, we suspect these differences in strategies had a major contribution to the time and error results for the Range task, and they are discussed in further detail in Section 6.6.

Looking at the time results further, we also notice that participants who had previous *VRUsage* and *TopoUsage* had different performance under some circumstances. Specifically, participants who had used VR previously had faster times in VR, and participants who had used topographic maps previously had faster times in both 3D modalities. The correlation between faster VR times and previous VR usage is understandable; VR usage among the general population is not yet ubiquitous, and performing tasks in VR does have a learning curve. We are, however, surprised by the correlation between previous topographic map usage and speed in the 3D modalities; we had expected to see the speed increase in the 2D modality instead. One possible explanation for this is that the topographic map usage skill transfers to a more general understanding of elevation data representations in any modality. We note that while effects of *TopoUsage* and *VRUsage* were found on the *time*, no such effects were found on *error* measure.

Finally, we see that participants perform faster with some datasets than with others. This is evident because the data originate from real elevation sources; thus we expect to see some variability between datasets. In addition to the elevation range and percentage of points that are touchable in the dataset (shown in Table 6.1), the random distribution of targets likely also plays a role in this. Because it is real data, some circle targets are inherently more difficult than others for the Advect and Range tasks, and some triplets of points for the Compare task are also more difficult than others (i.e., if two far away points have similar elevations). Though there are effects of *Dataset* on the *time* measure, no such effects were found on the *error* measures.

Confidence

Overall, participants' confidence ratings were fairly similar between modalities across the board, with two notable exceptions.

For the Advect task, participants were clearly most confident in 3D; the VR and Physical modalities both showed higher confidence rating than the 2D modality. From the qualitative feedback, we understand that in 2D many participants had difficulty interpreting the contour lines to follow the path downhill. While a majority of participants had prior experience with topographic maps, this was still a major source of trouble, and thus was a likely contributing factor to decreased confidence in 2D for the Advect task. We also attribute the differences in participants' confidence in 2D vs. 3D to the perceptual advantages that 3D visualizations provide when solving 3D tasks; these advantages are discussed further in Section 6.5.3.

For the Range task, participants were most confident in the 2D modality. In 2D, many participants used the strategy of counting contour lines and then looking up the numeric elevation, and participants preferred the certainty of the labeled contour lines to the uncertainty of performing a lookup based on the legend for the VR and Physical modalities. This difference in strategies was corroborated by the qualitative feedback, and likely influenced participants' confidence in the 2D modality for the Range task. We attribute these difference in confidence primarily to the visualization design of the 2D modality; this is discussed further in Section 6.6.

Comparison with Non-Spatial Studies

In comparison with prior studies comparing physical and digital modalities with non-spatial data, we observe a key similarity that this study echos the key results from those previous studies: *the Physical modality is at least as fast as digital modalities with comparable visualization design and interaction.* [Jansen et al., 2013, Danyluk et al., 2020].

We observe that somewhat surprisingly, the amount of errors committed on this study with spatial data was generally similar to the prior studies with non-spatial data; we use the Range and Compare task as points of comparison since they are common among the prior studies and our study. Both Jansen et al. and Danyluk et al. reported error rates for the Range task between 2% - 9% and Compare task between 10% - 14%, and opted not to report error difference between modalities or perform analysis on the errors. We saw errors between 3% - 11% for the Range task and 8% - 16% for the Compare task, roughly in the same range as the previous studies. This is interesting because with the spatial data in this study, participants were required to locate points and areas on a continuous surface as opposed to performing lookups based on discrete data.

Though we kept the Compare and Range tasks as similar as possible to prior studies (i.e., number of points on the grid), we see differences in the average time it took participants to complete the Compare and Range tasks in the Physical modality. For the Range task, Jansen et al. reported an average time of about 23 seconds, Danyluk et al. reported an average time of about 33 seconds, whereas ours was about 46 seconds. We attribute this study's much higher average time to the legend lookups necessary for the Physical and VR modalities. For the Compare task, Jansen et al. reported an average time of about 41 seconds, Danyluk et al. reported an average time of about 52 seconds, whereas ours was only about 36 seconds. We expect to see differences like this because some tasks are inherently different with spatial data; for the tasks in our study, participants were required to complete the tasks over *continuously-sampled* points and areas, rather than discrete points and areas used by prior studies.

6.5.2 Implications of Physicalization on Spatial Data Analysis

Beyond the main hypotheses, we also investigate the role of touch and the benefits and drawbacks of using data physicalizations for spatial data analysis tasks. We find that, unsurprisingly, the benefits of touch in performing data analysis tasks identified by prior studies are also present in this study on spatial data, perhaps even to a larger degree.

Role of Touch

From the results and feedback, we conclude that haptic perception of data can be especially helpful when interpreting the minute surface differences apparent in continuously sampled spatial data like those used in this study. While it is well-known that using the haptic channel enables a rich transfer of information between the data and the user, including material properties like roughness and friction [Hornecker et al., 2023], haptic perception is generally less accurate than visual perception for interpreting unfamiliar objects and data, and that using haptic perception alone forces a user to interpret data sequentially over time [McCormack et al., 2018]. Even so, a majority of participants mentioned that *touch* and *manipulation* were strategies they used to solve the spatial data analysis tasks. In fact, multiple participants reported strategies that relied strongly on touch: P8 stated their strategy was "*...to rely on my sense of touch to find the lowest and highest point and trace the path down the hill*", P11 stated that they would “*[p]hysically touch points of interest when eyes could not tell the difference*”, and P20 “*placed my fingers on the model and compared the touch sensation when required*”. Thus, touch plays a critical role in the comprehensibility and embodiment of physicalizations, especially for spatial data.

Participants touched the physicalization during more than 90% of Compare task trials, over two thirds of Advect trials, and over half of Range trials, and every participant touched the physicalization at least once during the study. So, even for the participants that didn't explicitly mention *touch* as a strategy, most implicitly touched the object to help them solve the tasks. Returning to the work on embodied cognition, one view is that physical objects help us to offload cognitive work into the environment [Wilson, 2002] rather than storing every bit of information in memory; this has been shown as an effective strategy for some visualization tasks [Jansen et al., 2013]. This offloading behavior was observed in a majority of trials, with participants using their fingers to mark points and "save" their locations, to trace paths downhill, and to draw lines between multiple points across the spatial data physicalization. We conclude that cognitive offloading and the affordance for touch that physical objects have [Klatzky and Peck, 2012] were the predominant factors leading most participants to touch the object, to find the most advantageous strategy to reduce their time and errors in solving the spatial data analysis tasks.

Though touch played an essential role in solving tasks, in the Physical modality most participants did not actually manipulate the physicalization by picking it up or rotating it; they were content to leave it on its base. One explanation for this is that participants were able to obtain a sufficient 3D understanding of the surface by the true, unambiguous depth cues provided by the physical object; it is also possible that the embodied interactions and haptic perception enabled by touching the object enabled most participants to complete the tasks to a satisfactory level (without moving it). We also suspect that the weight of the physicalizations may have been a contributing factor; though this wasn't reported in any of the qualitative feedback, participants may have felt fatigued by picking up the data object. Reviewing the time and error results with this in mind, we observe that this is another key result echoing earlier studies: the advantages of physicalization do not necessarily lie in its affordance for direct manipulation, rather it is the *touch* itself that matters.

Besides touch in the Physical modality, we notice that during almost a quarter of trials participants also touched the screen for the 2D modality (especially for the Compare task); rather than using the pan and zoom interactions with the mouse, some participants preferred to mark points with their fingers and then make comparisons. To illustrate this 2D strategy further, P15 reported that “*[f]or the compare task, I* *would locate the crosshairs of each of the three specified circles. I would then place a finger on each of those points so that it was easier to identify and thus compare the specified circles. After locating these circles, I similarly followed the process of looking at the color gradients, then continued to narrow down my answer by looking at the labelled contour lines.*" We note that the error for the 2D condition was the lowest in the Compare task, and we attribute at least some of this to participants using their embodied reference frame to offload some cognitive work into the environment. Additionally, multiple participants mentioned marking points with the 3D cursors in the VR modality, but since the virtual visualization was "popping out" of the screen, it was not possible to touch any part of the visualization.

Benefits and Drawbacks of Physicalization on Task Completion

The results from this study substantiate the theorized benefits that physicalizations provide, from true depth perception to bringing data into our embodied reference frame. A majority of participants mentioned the importance of the physicalization itself as a positive factor for their spatial data analysis due to the ability to physically touch and move the data object; P8 stated that "*[t]he ability to physically manipulate and feel the model helped in these tasks*", P24 stated "*[b]eing able to hold and move the map in 3D space helped be line up the elevations with the legend*", and P5 mentioned that the physicalization helped specifically for the Advect task: "*Having a tangible object to touch made the tracing/intersection task easier*". Beyond the direct touch advantages, others reported that the visual perception and manipulation of the physical object assisted P3; "*I can manipulate the physical object easily and compare [its] height when looking at the side of it*", P6 saying that "*...looking at the point from three views (front, left and top) helped identify the altitude*", and P17 noting that "*having the physical model let me see the perspective and angle exactly how [I] wanted*".

Participants also brought up some disadvantages to the Physical modality, namely its lack of interactivity and labeling. P5 mentioned that the "*lack of topography labels*" were problematic, and P15 similarly stated that it was "*a definite drawback to not have the contour line labels*". We also note that multiple participants used the strategy of "clipping" the legend through the virtual visualization in VR, and were frustrated that the Physical legend could not behave similarly because it was solid. Participants also thought that interactive annotation tools would be helpful for all three modalities, with P6 stating that "*it would be great if there could be some pin points when clicking on the selected cross marks*". The drawbacks of the Physical modality in terms of its interactivity are well-known in the literature [Djavaherpour et al., 2021], and are discussed in further detail in Section 6.6.

Finally, we observe that the Physical modality had higher ratio of positive feedback to negative feedback than either of the digital modalities. For 2D, participants identified 40 positive and 38 negative (ratio 1.05:1), for VR, there were 36 positive and 36 negative (ratio 1:1), and for Physical there were 38 positive and 18 negative (ratio 2.1:1). This result underscores the benefits that physicalizations provide in terms of their immediate relatability: since we interact with physical objects in the real world, we are intimately familiar with their affordances for touch and manipulation.

6.5.3 General Discussion

Thus far, the discussion has centered on the Physical modality and the implicit interactions it enables through touch and embodiment. However, there are other important results not directly tied to physicalizations, and these are discussed in the following paragraphs.

Unsurprisingly, the results indicate that 3D visualizations perform better than 2D displays for tasks that require a deep understanding of the 3D spatial structure of a dataset or object. This has been demonstrated numerous times in the literature, including navigation and positioning tasks [Tory et al., 2006], and spatial memory tasks [Cockburn and McKenzie, 2002]. In our results, the advantages of 3D are especially apparent in the Advect task; participants took less time in both 3D modalities than in 2D, and they reported higher confidence. This task required participants to form a mental model not only for the surface geometry but also the *gradient* of the geometry in order to solve the task; thus, 3D representations of the inherently 3D surface were helpful to completing this task.

Qualitative feedback from a majority of participants indicated that they found the 2D visualization *slow*, with P5 reporting that "*[i]t took some time for it to re-render the map whenever I changed zoom levels or panned the viewport...*", and P24 remarking that "*[s]ometimes the Letters and Numbers on the sides of the grids didn't render/load right away so I would have to wait a few seconds for those to pop up before I could figure out where to go on the map.*" This was not an intentional design decision and went unnoticed during testing and the pilot studies, so we conducted an investigation into the matter after the study was complete. We found from the video footage that participants panned and zoomed the 2D visualization about 5 times on average per trial, and using a screen recording we found that each refresh took on average 300ms. This equates to adding 1.5 seconds of loading on average per trial in the 2D modality.

6.6 Limitations and Future Work

6.6.1 Limitations

While this study introduces the first evaluation comparing spatial data physicalizations with digital visualizations, we note that it has some limitations. To begin, despite careful attention to the visual channels and usage of best practices for a GIS visualization, the 2D visualization design may have caused a confound. The contour lines in the 2D modality had 5 sub-contours, which enabled participants to effectively determine the elevations of points within about 40 meters, whereas the legends provided in both 3D modalities only provided steps of 200 meters. Additionally, both a colormap and contour lines were used in the 2D modality, enabling participants to use color to find high and low points. To avoid this confound, it would likely be better to use a hill shade instead of a colormap [Slocum et al., 2022] and to provide contour lines in the 3D modalities, not just in 2D. We expect that this design contributed to the reduced error rate in 2D, especially in the Range and Compare tasks.

Another potential factor that most likely contributed to the performance differences in the Range task was each modality's design for finding the numeric elevation at a given point on the visualization. For the 2D modality, numeric elevation labels were provided on the contour lines, which enabled participants to directly look up the minimum and maximum elevation. Conversely, for both 3D modalities, participants were required to look up elevations via a legend, a multi-step process that takes more time than a single numeric lookup on the 2D map. This was especially prevalent in VR because participants were not only required to look up the elevation but also needed to manually align the legend with the virtual visualization. Thus, for the Range task where these numeric lookups were necessary, we expect that the design for finding numeric elevations contributed to both reduced time and reduced error for the 2D modality and the increased time and error in both 3D modalities. So, while we can say for certain the Physical modality performs faster and with lower errors than digital modalities on the Advect task, we expect that we would see different results for the Compare and Range tasks if the legend and contour lines design considerations were addressed.

Following up on the qualitative results and feedback from participants, we reflect that there are likely further benefits to the Physical modality that our study was not able to capture. First, self-reported subjective measures like confidence are often questioned as a valid measure for a visualization's effectiveness; this is reflected in the study's results where confidence ratings were substantially the same except in two cases. Second, prior literature indicates that when investigating the value of a visualization beyond its performance and accuracy metrics, alternative measures may be required [Saket et al., 2016]. For example, though our qualitative feedback questionnaire at the end of each block provided some insight into participants' positive and negative sentiments about each modality and strategies they used, other methods like reaction cards [Merčun, 2014], or semi-structured interviews [Carpendale, 2008] with more space for open-ended responses may have been more effective. Finally, we note that evaluating physicalizations with respect to digital visualizations presents a unique set of opportunities and challenges. The perceptual, emotional, and embodied benefits of physicalizations are clear from the literature [Hornecker et al., 2023], but methods for effectively comparing these aspects of physicalizations with digital visualizations are less clear. For instance, information on individuals' differences in the "need for touch" [Peck and Childers, 2003] might be valuable to collect during recruitment, and further methods of modeling and evaluating emotional engagement with data are needed [Wang et al., 2019].

6.6.2 Future Work

In addition to the limitations of the current study, we highlight some areas where future work is needed.

The DEM data used in this study are a common type of spatial data and have been a target of physicalization for many years (e.g., [Schmitz, 2004, Tateosian et al., 2010, Millar et al., 2018]). However, these data – which are 2.5D and only encode one variable – only scratch the surface of the data that is possible and necessary for today's scientific sensemaking processes. There are even more interesting research questions and visualization tasks that one can ask about 3D multivariate volumetric spatial data, such as finding an area of high concentration in a particular variable, tracing the path a particle would take in a flow field, and selecting cutting planes in a volume to analyze data across two axes, and other classic examples of scientific visualization queries [Nielson et al., 1997]. Scientific data also frequently encompass multiple points in time. Time-varying data can be prohibitive to physicalize if the objective is to display time steps with animation. However, other approaches to time-varying visualization exist that can be more effective in some situations (e.g., interactive combined plots [Schroeder et al., 2012], summary statistics combined with cutting planes [Schroeder et al., 2014], and motion tracers [Keefe et al., 2009]). These alternative approaches have even been applied to physicalizations, such as ChronoFab [Kazi et al., 2016]. Despite the alternative approaches, animation remains an important feature to support for time-varying visualizations.

Usually, this animation is supported by digital interactivity. In this study we intentionally limited the interactivity to basic view manipulation operations to minimize the differences between modalities. However, we note that interactivity like zoom, filter, and details-on-demand (i.e., the information-seeking mantra [Shneiderman, 1996]) is a powerful and necessary tool for exploring data and finding the trends and relationships essential to comprehending the physical phenomena studied by today's scientific sensemaking processes. Multiple participants expressed a desire for such additional interactivity with the data (e.g., labels, filtering, and annotation), and adding these interactive elements has been considered best practice for making a visualization more comprehensible for decades [Upson et al., 1989]. Given the other work in this dissertation like Chapter 5's approach to touch sensing and digital-physical interactivity that enables multi-touch queries into 3D volumetric time-varying multivariate data, we see a future where such interactivity is inherent to *all* visualization modalities including data physicalizations. Thus, a critical step for future work is to create, apply, and evaluate interactive data exploration tools for both physical and digital visualizations, and to investigate the role that the embodied interactions enabled by physical data representations might have on the perceptibility, comprehensibility, and relatability of 3D multivariate spatial data.

6.7 Conclusion

This chapter introduced the first user study evaluation comparing the effectiveness of data physicalizations with digital visualizations for spatial data analysis tasks. The study reimagined prior visualization evaluation designs from the physicalization and scientific visualization literature to evaluate spatial data physicalizations compared with digital 2D GIS visualizations and VR visualizations. Time and error results were reported in addition to subjective measures of confidence and participant qualitative feedback for each visualization modality. The results of the study suggest that using data physicalizations for some spatial data analysis tasks can lead to substantial performance benefits, and that using a data physicalization provides benefits beyond performance like direct touch on the data, haptic perception, and leveraging our own human embodied reference frame to make data more comprehensible. Given these results, we believe that data physicalizations are worth considering as a means for spatial data analysis.

chapter

seven

Conclusion

7.1 Review of Contributions Towards Thesis Statement 123

7.2 Future Research 125

7.2.1 Track 1: Design and Perception of Data Physicalizations . . 125

7.2.2 Track 2: Integrating Physicalizations to Scientific Sensemaking Processes 125

7.2.3 Track 3: Adding Interactivity to Physicalizations 126

7.3 General Conclusions 126

7.1 Review of Contributions Towards Thesis Statement

In this chapter, I review the contributions of this dissertation towards the thesis statement, and identify important future research directions for the visualization and physicalization research communities.

The thesis statement of this dissertation is:

> *New approaches to extended reality visualization that encode data using physical rendering processes enable better comprehension of and more embodied relationships with spatial data by leveraging the perceptual benefits and tactile interactions that are missing from current digital visualizations.*

Each chapter has contributed a piece towards the greater whole of making extended reality data visualizations more graspable – both more comprehensible, and more embodied – using physical rendering processes.

Chapter 3 presented **a systematic design exploration of the first 3D printed visualizations encoding scalar data as glyphs on a 3D surface.** The work consisted of a series of design experiments towards a first-of-its-kind data physicalization, which combined a physical surface geometry with a glyph-based scalar data overlay. This work demonstrates an important addition to data physicalization research for two reasons. First, physicalizations typically are created from non-spatial data like statistical data, and thus do not transfer well to the needs of current scientific sensemaking processes. Second, of the data physicalizations that *do* represent spatial data, they only rarely include more than just the surface data. Our work addresses these concerns by creating a design for glyph-based spatial data physicalizations informed by current scientific visualization techniques and data physicalization guidelines. This work uses physical rendering processes to bring these multivariate visualizations into the physical world to provide more embodied interactions with the data.

Chapter 4 presented **a software architecture and user interface enabling the design of multivariate 3D spatial data visualizations using handcrafted physical media.** This work introduced a new way of interpreting physical rendering processes by using elements from the physical world to encode data in digital visualizations. The architecture and design user interface contributed by this chapter are significant to visualization research because they make it possible to engage the expertise of artists to effectively design multivariate 3D spatial data visualizations. With artists in the loop for visualization design, visualizations are made not only more expressive and engaging by leveraging the embodiment of physical objects, but can also depict more data in a comprehensible manner, a key need of current scientific sensemaking processes.

Chapter 5 presented **a new approach to querying 3D spatial data using multi-touch input directly on a data physicalization.** This work developed a new approach to using multi-touch input on an augmented data physicalization to make real-time queries into multivariate 3D spatial datasets. This work is significant to visualization and data physicalization research because it presents a new approach to interactive multi-touch sensing on a physicalization by first segmenting the physicalization into columns, then using the columns to transfer the force of the user's touch down to a touch sensitive tablet. While this work is still in its early stages, its performance evaluation is promising and domain science collaborators are excited about its potential to help them better comprehend their data through the embodied, tactile interactions enabled by making data more graspable.

Finally, Chapter 6 presented **an empirical study on the effectiveness of data physicalizations for spatial data analysis tasks when compared with state-of-the-art virtual reality and 2D geographic information science visualizations.** This research is significant because it is the first example of a user study comparing visualization modalities for *spatial* data analysis tasks. The design itself is novel, taking inspiration from both prior data physicalization studies and scientific visualization studies. In the study results, we find that unsurprisingly, data physicalizations outperform digital visualizations in many cases, and embodied interactions such as touching the physicalization are likely a contributing factor to this.

In summary, each of these individual explorations on physical rendering processes have the potential to make data more graspable: Reinterpreting scientific visualization and data physicalization design guidelines makes the types of data used by today's scientific sensemaking processes more graspable by bringing it into the physical world. Reimagining the design of digital 3D multivariate visualizations to leverage physical elements from the real world has the potential to make them more expressive and engaging, and enables a new level of embodiment for these digital visualizations. Reconstructing data physicalizations to support the digital, interactive queries necessary for today's scientific sensemaking processes makes data physicalizations more comprehensible. And, reassessing the strategies for evaluating data physicalizations can lead to deeper, more meaningful results when comparing physicalizations with traditional digital visualizations. Investigating the new ways that physical rendering processes can be used to make data more graspable is the motivation behind each of these works that contribute towards the greater whole of making data more comprehensible and more embodied for the benefit of all.

7.2 Future Research

Beyond the work in this dissertation, I identify three key tracks for future research regarding the use of data physicalizations for scientific sensemaking processes.

7.2.1 Track 1: Design and Perception of Data Physicalizations

While a few individual examples of evaluations centered on the design decisions for physicalizations exist, there is still a lack of comprehensive evaluation across the data physicalization community. Specifically, as identified by a formative article on data physicalization, more evaluation on the design and perception of data physicalizations is needed [Jansen et al., 2015]. Interesting research questions along this track might include:

1. What are the most effective physical marks to use to encode data for various visual and haptic channels?
2. How does the size of a spatial data physicalization affect the perception of the data?
3. How well do current assumptions underlying many visualization designs transfer to physicalizations?
4. How might the additional benefits that physicalization provides (e.g., diverse material properties, haptic perception, embodied interactions, and emotional connection) be evaluated?

7.2.2 Track 2: Integrating Physicalizations to Scientific Sensemaking Processes

As described earlier in this dissertation, the data physicalization community has been largely focused on non-spatial data. However, these data are not as applicable to scientific sensemaking processes, and though our work in Chapter 6 provides a first step in this direction, additional prototypes and experiments are needed that establish the utility of data physicalizations for scientific sensemaking processes. Promising research objectives along this track include:

1. Extending current digital fabrication technologies for use in scientific workflows; specifically, multi-color and multi-material manufacturing to facilitate the physicalization of multivariate and high dimensional data.

2. Conducting further experiments on the most effective ways to depict multivariate spatial data in a physical modality; for example, color, glyphs, physical texture, softness, squishiness, or tensile strength.

3. Researching new domain science application areas and seeking out collaborations that would benefit from physicalizing data.

7.2.3 Track 3: Adding Interactivity to Physicalizations

Another consideration identified by this dissertation and others (e.g., [Jansen et al., 2015, Djavaherpour et al., 2021]) is the current lack of interactivity in data physicalizations; this is necessary to support current scientific sensemaking processes which require quick changes and feedback to data queries and even steering a simulation while it's running. One approach to this is "augmented" data physicalizations like the one described in Chapter 5, but they currently are limited; they generally give strong preference to the digital visualization (e.g., augmented reality scatter plots [Bach et al., 2018]), or a preference to the physical at the expense of interactivity (e.g., [Khot et al., 2020]). Next-generation augmented physicalizations will push the boundaries between the digital and physical world, striving for a seamless data experience providing a "best-of-both-worlds" between digital interactivity and physical embodiment. Future research to enable such augmented physicalizations includes:

1. Developing new tracking and sensing technologies for interacting with data physicalizations in rich ways (e.g., gesture, multi-touch, voice input).

2. Comparing the performance of interactive visualizations across data physicalizations and digital visualization modalities.

3. Reconciling current research into active, robotically actuated data physicalizations with current scientific sensemaking processes.

4. Examining the diverse roles physicalizations play in sensemaking processes (e.g., a primary data object of interest, a contextual display, or a tangible interface to make queries into the data).

7.3 General Conclusions

The work presented in this dissertation explores how physical rendering processes can make extended reality visualizations more graspable.

Towards the comprehensibility dimension of graspability, I contribute the first study comparing the effectiveness of data physicalizations for spatial data analysis tasks. From the results of this study, I conclude that data representations created through physical rendering processes can not only perform faster with fewer errors than their digital counterparts for certain spatial data analysis tasks, but the act of *touch* on these physicalizations is valuable for its haptic perceptual benefits and the embodied interactions it provides. I also contribute a new approach to enabling spatial data queries necessary for today's scientific sensemaking processes, by combining multi-touch sensing on a data physicalization with immersive augmented reality. From these explorations on interactive querying and the comprehensibility benefits of physicalization shown by the study, I conclude that there is significant potential for data physicalizations with digital interactivity to be used in scientific sensemaking processes, especially in terms of the possibilities for collaboration and engagement that they provide.

Towards the embodiment dimension of graspability, I contribute multiple examples of data physicalizations representative of the 3D multivariate spatial data necessitated by today's scientific sensemaking processes, including a methodical exploration on the design of data physicalizations encoding 3D geometry and scalar variables as glyphs on 3D printed surfaces. Additionally, I contribute a visualization design user interface and architecture that enables visualization designers to use physical artifacts from the real world to encode variables in a 3D multivariate visualization. In this work, I conclude that because of the physical origin of the data encodings, some of the emotional, relatable, and embodied benefits of physical data representations are present through *imagined touch*, even though the these results are displayed in a digital extended reality visualization.

In conclusion, extended reality data visualizations created through physical sensemaking processes convey a range of information through their geometric and material properties, and their capacity to support tactile exploration, haptic perception, and embodied interactions enables us to develop a better comprehension of the universe. Ultimately, the work presented in this dissertation anticipates the future of extended reality visualizations where the boundaries between the digital and physical are blurred: where scientists can directly touch a data physicalization to make queries into their 3D multivariate spatial datasets, and query results are displayed as digital extended reality visualizations blended with the physicalization; where artists collaborate with scientists to make the most comprehensible and engaging visualizations using materials directly from their studios; and where the benefits of *both* highly interactive digital visualizations and embodied physical visualizations are shared by all.

Bibliography

[Abawi et al., 2004] Abawi, D. F., Bienwald, J., and Dorner, R. (2004). Accuracy in optical tracking with fiducial markers: An accuracy function for ARToolKit. In *Third IEEE and ACM International Symposium on Mixed and Augmented Reality*, pages 260–261. IEEE.

[Abram et al., 2021] Abram, G., Samsel, F., Petersen, M. R., Asay-Davis, X., Comeau, D., Price, S. F., and Potel, M. (2021). Antarctic water masses and ice shelves: Visualizing the physics. *IEEE Computer Graphics and Applications*, 41(1):35–41.

[Acevedo et al., 2001] Acevedo, D., Vote, E., Laidlaw, D. H., and Joukowsky, M. S. (2001). Archaeological data visualization in VR: Analysis of lamp finds at the Great Temple of Petra, a case study. In *Proceedings Visualization, 2001. VIS'01.*, pages 493–597. IEEE.

[Ahrens et al., 2005] Ahrens, J., Geveci, B., and Law, C. (2005). Paraview: An end-user tool for large data visualization. *The visualization handbook*, 717(8).

[Alonso et al., 2018] Alonso, L., Zhang, Y. R., Grignard, A., Noyman, A., Sakai, Y., ElKatsha, M., Doorley, R., and Larson, K. (2018). Cityscope: A data-driven interactive simulation tool for urban design. Use case volpe. In *International Conference on Complex Systems*, pages 253–261. Springer.

[Ang et al., 2019] Ang, K. D., Samavati, F. F., Sabokrohiyeh, S., Garcia, J., and Elbaz, M. S. (2019). Physicalizing cardiac blood flow data via 3D printing. *Computers and Graphics*, 85(C):42–54.

[Bach et al., 2018] Bach, B., Sicat, R., Beyer, J., Cordeil, M., and Pfister, H. (2018). The Hologram in My Hand: How Effective is Interactive Exploration of 3D Visualizations in Immersive Tangible Augmented Reality? *IEEE Transactions on Visualization and Computer Graphics*, 24(1):457–467.

[Bader et al., 2018] Bader, C., Kolb, D., Weaver, J. C., Sharma, S., Hosny, A., Costa, J., and Oxman, N. (2018). Making data matter: Voxel printing for the digital fabrication of data across scales and domains. *Science Advances*, 4(5):eaas8652.

[Bae et al., 2022] Bae, S. S., Zheng, C., West, M. E., Do, E. Y.-L., Huron, S., and Szafir, D. A. (2022). Making Data Tangible: A Cross-disciplinary Design Space for Data Physicalization.

[Bailey et al., 1998] Bailey, M. J., Schulten, K., and Johnson, J. E. (1998). The use of solid physical models for the study of macromolecular assembly. page 7.

[Berard and Louis, 2017] Berard, F. and Louis, T. (2017). The Object Inside: Assessing 3D Examination with a Spherical Handheld Perspective-Corrected Display. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pages 4396–4404, New York, NY, USA. Association for Computing Machinery.

[Bertini et al., 2008] Bertini, E., Plaisant, C., Perer, A., and Santucci, G. (2008). BELIV'08: BEyond time and errors— novel evaLuation methods for Information Visualization.

[Best, 1996] Best, S. (1996). Perceptual and oculomotor implications of interpupillary distance settings on a head-mounted virtual display. In *Proceedings of the IEEE 1996 National Aerospace and Electronics Conference NAECON 1996*, volume 1, pages 429–434, Dayton, OH, USA. IEEE.

[Bilgili, 2017] Bilgili, D. (2017). *A Data Physicalization Pipeline Enhanced with Augmented Reality*. PhD thesis.

[Billinghurst et al., 2008] Billinghurst, M., Kato, H., Poupyrev, I., et al. (2008). Tangible augmented reality. *Acm siggraph asia*, 7(2):1–10.

[Bobrich and Otto, 2002] Bobrich, J. and Otto, S. (2002). Augmented maps. *International Archives of Photogrammetry Remote Sensing and Spatial Information Sciences*, 34(4):502–505.

[Boegehold, 1963] Boegehold, A. L. (1963). Toward a Study of Athenian Voting Procedure. *Hesperia*, 32(4):366.

[Bonis et al., 2022] Bonis, M. D., Nguyen, H., and Bourdot, P. (2022). A Literature Review of User Studies in Extended Reality Applications for Archaeology. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 92–101, Singapore, Singapore. IEEE.

[Boucheny et al., 2009] Boucheny, C., Bonneau, G.-P., Droulez, J., Thibault, G., and Ploix, S. (2009). A perceptive evaluation of volume rendering techniques. *ACM Transactions on Applied Perception*, 5(4):1–24.

[Brehmer and Munzner, 2013] Brehmer, M. and Munzner, T. (2013). A Multi-Level Typology of Abstract Visualization Tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2376–2385.

[Bridson, 2007] Bridson, R. (2007). Fast Poisson disk sampling in arbitrary dimensions. In *ACM SIGGRAPH 2007 Sketches*, page 22, San Diego California. ACM.

[Bruder et al., 2011] Bruder, G., Steinicke, F., Wieland, P., and Lappe, M. (2011). Tuning self-motion perception in virtual reality with visual illusions. *IEEE Transactions on Visualization and Computer Graphics*, 18(7):1068–1078.

[Bryson, 1992] Bryson, S. (1992). Virtual Reality Takes on Real Physics Applications. *Computers in Physics*, 6(4):346.

[Bryson, 1996] Bryson, S. (1996). Virtual reality in scientific visualization. *Communications of the ACM*, 39(5):62–71.

[Büschel et al., 2018] Büschel, W., Chen, J., Dachselt, R., Drucker, S., Dwyer, T., Görg, C., Isenberg, T., Kerren, A., North, C., and Stuerzlinger, W. (2018). Interaction for Immersive Analytics. In Marriott, K., Schreiber, F., Dwyer, T., Klein, K., Riche, N. H., Itoh, T., Stuerzlinger, W., and Thomas, B. H., editors, *Immersive Analytics*, volume 11190, pages 95–138. Springer International Publishing, Cham.

[Buxton, 1997] Buxton, B. (1997). Artists and the art of the luthier. *ACM SIGGRAPH Computer Graphics*, 31(1):10–11.

[Buxton, 2010] Buxton, B. (2010). *Sketching User Experiences: Getting the Design Right and the Right Design*. Morgan kaufmann.

[Buxton et al., 1985] Buxton, W., Hill, R., and Rowley, P. (1985). Issues and techniques in touch-sensitive tablet input. In *Proceedings of the 12th Annual Conference on Computer Graphics and Interactive Techniques*, pages 215–224.

[Carpendale, 2008] Carpendale, S. (2008). Evaluating Information Visualizations. In Kerren, A., Stasko, J. T., Fekete, J.-D., and North, C., editors, *Information Visualization*, volume 4950, pages 19–45. Springer Berlin Heidelberg, Berlin, Heidelberg.

[Carroll and Blauch, 2017] Carroll, F. A. and Blauch, D. N. (2017). 3D printing of molecular models with calculated geometries and p orbital isosurfaces. *Journal of Chemical Education*, 94(7):886–891.

[Chang et al., 2017] Chang, J. S.-K., Yeboah, G., Doucette, A., Clifton, P., Nitsche, M., Welsh, T., and Mazalek, A. (2017). TASC: Combining Virtual Reality with Tangible and Embodied Interactions to Support Spatial Cognition. In *Proceedings of the 2017 Conference on Designing Interactive Systems*, pages 1239–1251, Edinburgh United Kingdom. ACM.

[Chen et al., 2020] Chen, Z., Su, Y., Wang, Y., Wang, Q., Qu, H., and Wu, Y. (2020). MARVisT: Authoring Glyph-Based Visualization in Mobile Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics*, 26(8):2645–2658.

[Childs et al., 2012] Childs, H., Brugger, E., Whitlock, B., Meredith, J., Ahern, S., Pugmire, D., Biagas, K., Miller, M., Harrison, C., Weber, G. H., Krishnan, H., Fogal, T., Sanderson, A., Garth, C., Bethel, E. W., Camp, D., Rübel, O., Durant, M., Favre, J. M., and Navrátil, P. (2012). VisIt: An end-user tool for visualizing and analyzing very large data. In *High Performance Visualization–Enabling Extreme-Scale Scientific Insight*, pages 357–372.

[Cockburn and McKenzie, 2002] Cockburn, A. and McKenzie, B. (2002). Evaluating the Effectiveness of Spatial Memory in 2D and 3D Physical and Virtual Environments. *Spatial Cognition*, (4):8.

[Coffey et al., 2012] Coffey, D., Korsakov, F., Ewert, M., Hagh-Shenas, H., Thorson, L., Ellingson, A., Nuckley, D., and Keefe, D. F. (2012). Visualizing motion data in virtual reality: Understanding the roles of animation, interaction, and static presentation. In *Computer Graphics Forum*, volume 31, pages 1215–1224. Wiley Online Library.

[Coffey et al., 2011a] Coffey, D., Malbraaten, N., Le, T., Borazjani, I., Sotiropoulos, F., and Keefe, D. F. (2011a). Slice WIM: A multi-surface, multi-touch interface for overview+ detail exploration of volume datasets in virtual reality. In *Symposium on Interactive 3D Graphics and Games*, pages 191–198.

[Coffey et al., 2011b] Coffey, D., Malbraaten, N., Le, T. B., Borazjani, I., Sotiropoulos, F., Erdman, A. G., and Keefe, D. F. (2011b). Interactive slice WIM: Navigating and interrogating volume data sets using a multisurface, multitouch VR interface. *IEEE Transactions on Visualization and Computer Graphics*, 18(10):1614–1626.

[Cohen, 2013] Cohen, J. (2013). *Statistical Power Analysis for the Behavioral Sciences*. Academic press.

[Çöltekin et al., 2020] Çöltekin, A., Lochhead, I., Madden, M., Christophe, S., Devaux, A., Pettit, C., Lock, O., Shukla, S., Herman, L., Stachoň, Z., Kubíček, P., Snopková, D., Bernardes, S., and Hedley, N. (2020). Extended Reality in Spatial Sciences: A Review of Research Challenges and Future Directions. *ISPRS International Journal of Geo-Information*, 9(7):439.

[Cousineau, 2019] Cousineau, D. (2019). Correlation-adjusted standard errors and confidence intervals for within-subject designs: A simple multiplicative approach. *The Quantitative Methods for Psychology*, 15(3):226–241.

[Cox, 1988] Cox, D. J. (1988). Using the supercomputer to visualize higher dimensions: An artist's contribution to scientific visualization. *Leonardo*, 21(3):233–242.

[Cox, 1991] Cox, D. J. (1991). Collaborations in art/science: Renaissance teams. *The Journal of biocommunication*, 18(2):15–24.

[Cruz-Neira et al., 1993] Cruz-Neira, C., Leigh, J., Papka, M., Barnes, C., Cohen, S. M., Das, S., Engelmann, R., Hudson, R., Roy, T., Siegel, L., et al. (1993). Scientists in wonderland: A report on visualization applications in the CAVE virtual reality environment. In *Proceedings of 1993 IEEE Research Properties in Virtual Reality Symposium*, pages 59–66.

[Dadkhahfard et al., 2018] Dadkhahfard, S., Etemad, K., Brosz, J., and Samavati, F. (2018). Area Preserving Dynamic Geospatial Visualization on Physical Globe:. In *Proceedings of the 13th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, pages 309–318, Funchal, Madeira, Portugal. SCITEPRESS - Science and Technology Publications.

[Danyluk et al., 2020] Danyluk, K., Ulusoy, T., Wei, W., and Willett, W. (2020). Touch and Beyond: Comparing Physical and Virtual Reality Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 28(4):1930–1940.

[Djavaherpour et al., 2021] Djavaherpour, H., Samavati, F., Mahdavi-Amiri, A., Yazdanbakhsh, F., Huron, S., Levy, R., Jansen, Y., and Oehlberg, L. (2021). Data to physicalization: A survey of the physical rendering process. *Computer Graphics Forum*, 40(3):569–598.

[Dragicevic et al., 2021] Dragicevic, P., Jansen, Y., and Vande Moere, A. (2021). Data Physicalization. In Vanderdonckt, J., Palanque, P., and Winckler, M., editors, *Handbook of Human Computer Interaction*, pages 1–51. Springer International Publishing, Cham.

[Drogemuller et al., 2021] Drogemuller, A., Cunningham, A., Walsh, J. A., Baumeister, J., Smith, R. T., and Thomas, B. H. (2021). Haptic and Visual Comprehension of a 2D Graph Layout Through Physicalisation. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–16, Yokohama Japan. ACM.

[Duffy, 2014] Duffy, C. (2014). Abyss Table - Scale Model of Deep Sea as Furniture. http://dataphys.org/list/abyss-table/.

[Dwyer, 2004] Dwyer, T. (2004). Two-and-a-Half-Dimensional Visualisation of Relational Networks. page 257.

[El Beheiry et al., 2019] El Beheiry, M., Doutreligne, S., Caporal, C., Ostertag, C., Dahan, M., and Masson, J.-B. (2019). Virtual Reality: Beyond Visualization. *Journal of Molecular Biology*, 431(7):1315–1321.

[Eynard and Jenny, 2016] Eynard, J. D. and Jenny, B. (2016). Illuminated and shadowed contour lines: Improving algorithms and evaluating effectiveness. *International Journal of Geographical Information Science*, pages 1–21.

[Fang et al., 2015] Fang, B., Guo, D., Sun, F., Liu, H., and Wu, Y. (2015). A robotic hand-arm teleoperation system using human arm/hand with a novel data glove. In *2015 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 2483–2488, Zhuhai. IEEE.

[Fang et al., 2017] Fang, H., Walton, S., Delahaye, E., Harris, J., Storchak, D. A., and Chen, M. (2017). Categorical Colormap Optimization with Visualization Case Studies. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):871–880.

[Farooq et al., 2016] Farooq, H., Xu, J., Nam, J. W., Keefe, D. F., Yacoub, E., Georgiou, T., and Lenglet, C. (2016). Microstructure Imaging of Crossing (MIX) White Matter Fibers from diffusion MRI. *Scientific Reports*, 6(1):38927.

[Faul et al., 2009] Faul, F., Erdfelder, E., Buchner, A., and Lang, A.-G. (2009). Statistical power analyses using G* Power 3.1: Tests for correlation and regression analyses. *Behavior research methods*, 41(4):1149–1160.

[Fette and Melnikov, 2011] Fette, I. and Melnikov, A. (2011). The websocket protocol.

[Fitzmaurice et al., 1995] Fitzmaurice, G. W., Ishii, H., and Buxton, W. A. S. (1995). Bricks: Laying the foundations for graspable user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '95*, pages 442–449, Denver, Colorado, United States. ACM Press.

[Follmer et al., 2013] Follmer, S., Leithinger, D., Olwal, A., Hogge, A., and Ishii, H. (2013). inFORM: Dynamic Physical Affordances and Constraints through Shape and Object Actuation. page 10.

[Fonnet and Prie, 2021] Fonnet, A. and Prie, Y. (2021). Survey of Immersive Analytics. *IEEE Transactions on Visualization and Computer Graphics*, 27(3):2101–2122.

[Forsberg et al., 2009] Forsberg, A., Jian Chen, and Laidlaw, D. (2009). Comparing 3D Vector Field Visualization Methods: A User Study. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1219–1226.

[Fuchs et al., 2017] Fuchs, J., Isenberg, P., Bezerianos, A., and Keim, D. (2017). A Systematic Review of Experimental Studies on Data Glyphs. *IEEE Transactions on Visualization and Computer Graphics*, 23(7):1863–1879.

[Gao et al., 2023] Gao, L., Irani, P., Subramanian, S., Prabhakar, G., Martinez Plasencia, D., and Hirayama, R. (2023). DataLev: Mid-air Data Physicalisation Using Acoustic Levitation. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–14, Hamburg Germany. ACM.

[Gibson, 2014] Gibson, J. J. (2014). *The Ecological Approach to Visual Perception: Classic Edition*. Psychology press.

[Gillet et al., 2004] Gillet, A., Sanner, M., Stoffler, D., Goodsell, D., and Olson, A. (2004). Augmented reality with tangible auto-fabricated models for molecular biology applications. In *IEEE Visualization 2004*, pages 235–241. IEEE.

[Gillet et al., 2005] Gillet, A., Sanner, M., Stoffler, D., and Olson, A. (2005). Tangible Interfaces for Structural Molecular Biology. *Structure*, 13(3):483–491.

[Greenberg et al., 2013] Greenberg, I., Xu, D., and Kumar, D. (2013). *Processing: Creative Coding and Generative Art in Processing 2*. Apress.

[Grossman, 2002] Grossman, B. (2002). Bathsheba Grossman's Crystal Engravings. http://dataphys.org/list/scientific-visualization-in-crystal/.

[Groves et al., 2019] Groves, L. A., Carnahan, P., Allen, D. R., Adam, R., Peters, T. M., and Chen, E. C. (2019). Accuracy assessment for the co-registration between optical and VIVE head-mounted display tracking. *International journal of computer assisted radiology and surgery*, 14(7):1207–1215.

[Gwilt, 2022] Gwilt, I., editor (2022). *Making Data: Materializing Digital Information*. Bloomsbury Publishing Plc, 1 edition.

[Hedges and Olkin, 2014] Hedges, L. V. and Olkin, I. (2014). *Statistical Methods for Meta-Analysis*. Academic press.

[Hedley et al., 2002] Hedley, N. R., Billinghurst, M., Postner, L., May, R., and Kato, H. (2002). Explorations in the use of augmented reality for geographic visualization. *Presence: Teleoperators & Virtual Environments*, 11(2):119–133.

[Herman and Keefe, 2018] Herman, B. and Keefe, D. F. (2018). Boxcars on Potatoes: Exploring the Design Language for Tangible Visualizations of Scalar Data Fields on 3D Surfaces. page 10.

[Herman et al., 2021] Herman, B., Omdal, M., Zeller, S., Richter, C. A., Samsel, F., Abram, G., and Keefe, D. F. (2021). Multi-Touch Querying on Data Physicalizations in Immersive AR. *Proceedings of the ACM on Human-Computer Interaction*, 5(ISS):1–20.

[Herman et al., 2020] Herman, B., Samsel, F., Bares, A., Johnson, S., Abram, G., and Keefe, D. F. (2020). Printmaking, Puzzles, and Studio Closets: Using Artistic Metaphors to Reimagine the User Interface for Designing Immersive Visualizations.

[Hogan and Hornecker, 2012] Hogan, T. and Hornecker, E. (2012). How Does Representation Modality Affect User-Experience of Data Artifacts? In Hutchison, D., Kanade, T., Kittler, J., Kleinberg, J. M., Mattern, F., Mitchell, J. C., Naor, M., Nierstrasz, O., Pandu Rangan, C., Steffen, B., Sudan, M., Terzopoulos, D., Tygar, D., Vardi, M. Y., Weikum, G., Magnusson, C., Szymczak, D., and Brewster, S., editors, *Haptic and Audio Interaction Design*, volume 7468, pages 141–151. Springer Berlin Heidelberg, Berlin, Heidelberg.

[Holstius et al., 2004] Holstius, D., Kembel, J., Hurst, A., Wan, P.-H., and Forlizzi, J. (2004). Infotropism: Living and robotic plants as interactive displays. In *Proceedings of the 5th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques*, pages 215–221, Cambridge MA USA. ACM.

[Hornecker, 2011] Hornecker, E. (2011). The role of physicality in tangible and embodied interactions. *Interactions*, 18(2):19–23.

[Hornecker and Buur, 2006] Hornecker, E. and Buur, J. (2006). Getting a Grip on Tangible Interaction: A Framework on Physical Space and Social Interaction.

[Hornecker et al., 2023] Hornecker, E., Hogan, T., Hinrichs, U., and Van Koningsbruggen, R. (2023). A Design Vocabulary for Data Physicalization. *ACM Transactions on Computer-Human Interaction*, page 3617366.

[Isenberg et al., 2008] Isenberg, T., Everts, M. H., Grubert, J., and Carpendale, S. (2008). Interactive exploratory visualization of 2D vector fields. In *Computer Graphics Forum*, volume 27, pages 983–990. Wiley Online Library.

[Ishii and Ullmer, 1997] Ishii, H. and Ullmer, B. (1997). Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, pages 234–241, Atlanta Georgia USA. ACM.

[Jackson et al., 2012a] Jackson, B., Coffey, D., and Keefe, D. F. (2012a). Force Brushes: Progressive data-driven haptic selection and filtering for multi-variate flow visualizations. *Proceedings of EuroVis 2012*, pages 7–11.

[Jackson et al., 2012b] Jackson, B., Coffey, D., Thorson, L., Schroeder, D., Ellingson, A. M., Nuckley, D. J., and Keefe, D. F. (2012b). Toward mixed method evaluations of scientific visualizations and design process as an evaluation tool. In *Proceedings of the 2012 BELIV Workshop: Beyond Time and Errors - Novel Evaluation Methods for Visualization*, pages 1–6, Seattle Washington USA. ACM.

[Jackson and Keefe, 2019] Jackson, B. and Keefe, D. F. (2019). From painting to widgets, 6-DOF and bimanual input beyond pointing.

[Jackson et al., 2013] Jackson, B., Lau, T. Y., Schroeder, D., Toussaint, K. C., and Keefe, D. F. (2013). A Lightweight Tangible 3D Interface for Interactive Visualization of Thin Fiber Structures. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2802–2809.

[Jackson et al., 2012c] Jackson, B., Schroeder, D., and Keefe, D. F. (2012c). Nailing down multi-touch: Anchored above the surface interaction for 3D modeling and navigation. In *Proceedings of Graphics Interface 2012*, pages 181–184. Citeseer.

[Jansen and Dragicevic, 2013] Jansen, Y. and Dragicevic, P. (2013). An Interaction Model for Visualizations Beyond The Desktop. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2396–2405.

[Jansen et al., 2013] Jansen, Y., Dragicevic, P., and Fekete, J.-D. (2013). Evaluating the efficiency of physical visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2593–2602, Paris France. ACM.

[Jansen et al., 2015] Jansen, Y., Dragicevic, P., Isenberg, P., Alexander, J., Karnik, A., Kildal, J., Subramanian, S., and Hornbæk, K. (2015). Opportunities and Challenges for Data Physicalization. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 3227–3236, Seoul Republic of Korea. ACM.

[Jansen and Hornbaek, 2016] Jansen, Y. and Hornbaek, K. (2016). A Psychophysical Investigation of Size as a Physical Variable. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):479–488.

[Jeong et al., 2020] Jeong, H., Asay-Davis, X. S., Turner, A. K., Comeau, D. S., et al. (2020). Impacts of ice-shelf melting on water-mass transformation in the southern ocean from E3SM simulations. *Journal of Climate*, 33(13):5787–5807.

[Johnson, 2020] Johnson, S. (2020). *Palpable Visualizations: Techniques for Creatively Designing Discernible and Accessible Visualizations Grounded in the Physical World*. PhD thesis, University of Minnesota.

[Johnson et al., 2019a] Johnson, S., Orban, D., Runesha, H. B., Meng, L., Juhnke, B., Erdman, A., Samsel, F., and Keefe, D. F. (2019a). Bento box: An interactive and zoomable small multiples technique for visualizing 4d simulation ensembles in virtual reality. *Frontiers in Robotics and AI*, 6:61.

[Johnson et al., 2019b] Johnson, S., Samsel, F., Abram, G., Olson, D., Solis, A. J., Herman, B., Wolfram, P. J., Lenglet, C., and Keefe, D. F. (2019b). Artifact-Based Rendering: Harnessing Natural and Traditional Visual Media for More Expressive and Engaging 3D Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1.

[Kahn, 1996] Kahn, N. (1996). Ned Kahn's Wind-Visualizing Facades. http://dataphys.org/list/ned-kahns-wind-visualizing-facades/.

[Kaplan, 1997] Kaplan, E. (1997). Anecdotes. *IEEE Annals of the History of Computing*, 19(2):62–69.

[Kaspersen et al., 2021] Kaspersen, M. H., Bilstrup, K.-E. K., and Petersen, M. G. (2021). The Machine Learning Machine: A Tangible User Interface for Teaching Machine Learning. In *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction*, pages 1–12, Salzburg Austria. ACM.

[Kazi et al., 2016] Kazi, R. H., Grossman, T., Mogk, C., Schmidt, R., and Fitzmaurice, G. (2016). ChronoFab: Fabricating Motion. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 908–918, San Jose California USA. ACM.

[Keefe et al., 2009] Keefe, D., Ewert, M., Ribarsky, W., and Chang, R. (2009). Interactive Coordinated Multiple-View Visualization of Biomechanical Motion Data. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1383–1390.

[Keefe et al., 2005] Keefe, D., Karelitz, D., Vote, E., and Laidlaw, D. (2005). Artistic collaboration in designing VR visualizations. *Computer Graphics and Applications, IEEE*, 25(2):18–23.

[Keefe et al., 2008a] Keefe, D. F., Acevedo, D., Miles, J., Drury, F., Swartz, S. M., and Laidlaw, D. H. (2008a). Scientific Sketching for Collaborative VR Visualization Design. *IEEE Transactions on Visualization and Computer Graphics*, 14(4):835–847.

[Keefe et al., 2022a] Keefe, D. F., Altheimer, R., Johnson, A. J., Mahmoudi, M., Moe, P., Rockcastle, M., Swackhamer, M., and Wittkamper, A. (2022a). Orbacles: MINN_LAB design collective. In *Making with Data*, pages 221–232. AK Peters/CRC Press.

[Keefe et al., 2022b] Keefe, D. F., Herman, B., Nam, J. W., Johnson, S., and Orban, D. (2022b). Hybrid Data Constructs: Interacting with Biomedical Data in Augmented Spaces. In *Making Data: Materializing Digital Information*, pages 169–182. Bloomsbury Publishing Plc, 1 edition.

[Keefe et al., 2020] Keefe, D. F., Samsel, F., and Herman, B. (2020). Artifact-Based Rendering: VR visualization by hand. IEEE VIS Tutorial, Virtual.

[Keefe et al., 2008b] Keefe, D. F., Zeleznik, R. C., and Laidlaw, D. H. (2008b). Technote: Dynamic dragging for input of 3D trajectories. In *2008 IEEE Symposium on 3D User Interfaces*, pages 51–54. IEEE.

[Kendrew et al., 1958] Kendrew, J. C., Bodo, G., Dintzis, H. M., Parrish, RG., Wyckoff, H., and Phillips, D. C. (1958). A three-dimensional model of the myoglobin molecule obtained by x-ray analysis. *Nature*, 181(4610):662–666.

[Khot et al., 2020] Khot, R. A., Hjorth, L., and Mueller, F. (2020). Shelfie: A Framework for Designing Material Representations of Physical Activity Data. *ACM Transactions on Computer-Human Interaction*, 27(3):1–52.

[Khot et al., 2014] Khot, R. A., Hjorth, L., and Mueller, F. F. (2014). Understanding physical activity through 3D printed material artifacts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3835–3844, Toronto Ontario Canada. ACM.

[Kim et al., 2003] Kim, S., Hagh-Shenas, H., and Interrante, V. (2003). Showing shape with texture: Two directions seem better than one. In Rogowitz, B. E. and Pappas, T. N., editors, *Electronic Imaging 2003*, page 332, Santa Clara, CA.

[Kindlmann and Westin, 2006] Kindlmann, G. and Westin, C.-f. (2006). Diffusion Tensor Visualization with Glyph Packing. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1329–1336.

[Kirshenbaum et al., 2020] Kirshenbaum, N., Hutchison, J., Theriot, R., Kobayashi, D., and Leigh, J. (2020). Data in Context: Engaging Audiences with 3D Physical Geo-Visualization. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–9, Honolulu HI USA. ACM.

[Klatzky and Peck, 2012] Klatzky, R. L. and Peck, J. (2012). Please Touch: Object Properties that Invite Touch. *IEEE Transactions on Haptics*, 5(2):139–147.

[Klein et al., 2012] Klein, T., Guéniat, F., Pastur, L., Vernier, F., and Isenberg, T. (2012). A design study of direct-touch interaction for exploratory 3D scientific visualization. In *Computer Graphics Forum*, volume 31, pages 1225–1234. Wiley Online Library.

[Kohli, 2010] Kohli, L. (2010). Redirected touching: Warping space to remap passive haptics. In *2010 IEEE Symposium on 3D User Interfaces (3DUI)*, pages 129–130. IEEE.

[Kosara et al., 2003] Kosara, R., Healey, C. G., Interrante, V., Laidlaw, D. H., and Ware, C. (2003). Thoughts on User Studies: Why, How, and When.

[Kramida, 2016] Kramida, G. (2016). Resolving the Vergence-Accommodation Conflict in Head-Mounted Displays. *IEEE Transactions on Visualization and Computer Graphics*, 22(7):1912–1931.

[Kriz, 2007] Kriz, R. D. (2007). Thermodynamic Case Study: Gibbs Thermodynamic Graphical Method. https://esm.rkriz.net/classes/ESM4714/methods/Gibbs.html.

[Laha et al., 2015] Laha, B., Bowman, D. A., Laidlaw, D. H., and Socha, J. J. (2015). A classification of user tasks in visual analysis of volume data. In *2015 IEEE Scientific Visualization Conference (SciVis)*, pages 1–8, Chicago, IL, USA. IEEE.

[Laidlaw et al., 2005] Laidlaw, D., Kirby, R., Jackson, C., Davidson, J., Miller, T., da Silva, M., Warren, W., and Tarr, M. (2005). Comparing 2D Vector Field Visualization Methods: A User Study. *IEEE Transactions on Visualization and Computer Graphics*, 11(01):59–70.

[Lakens et al., 2018] Lakens, D., Scheel, A. M., and Isager, P. M. (2018). Equivalence Testing for Psychological Research: A Tutorial. *Advances in Methods and Practices in Psychological Science*, 1(2):259–269.

[Laramee et al., 2004a] Laramee, R., Weiskopf, D., Schneider, J., and Hauser, H. (2004a). Investigating swirl and tumble flow with a comparison of visualization techniques. In *IEEE Visualization 2004*, pages 51–58, Austin, TX, USA. IEEE Comput. Soc.

[Laramee et al., 2004b] Laramee, R. S., Hauser, H., Doleisch, H., Vrolijk, B., Post, F. H., and Weiskopf, D. (2004b). The State of the Art in Flow Visualization: Dense and Texture-Based Techniques. *Computer Graphics Forum*, 23(2):203–221.

[LaViola and Zeleznik, 1999] LaViola, J. and Zeleznik, R. (1999). Flex and pinch: A case study of whole hand input design for virtual environment interaction. In *Proceedings of the Second IASTED International Conference on Computer Graphics and Imaging*, pages 221–225.

[Laycock and Day, 2007] Laycock, S. and Day, A. (2007). A Survey of Haptic Rendering Techniques. *Computer Graphics Forum*, 26(1):50–65.

[Le et al., 2015] Le, A. T., Cole, G. G., and Wilkins, A. J. (2015). Assessment of trypophobia and an analysis of its visual precipitation. *Quarterly Journal of Experimental Psychology*, 68(11):2304–2322.

[Le Goc et al., 2019] Le Goc, M., Perin, C., Follmer, S., Fekete, J.-D., and Dragicevic, P. (2019). Dynamic Composite Data Physicalization Using Wheeled Micro-Robots. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):737–747.

[Lederman and Klatzky, 2009] Lederman, S. J. and Klatzky, R. L. (2009). Haptic perception: A tutorial. *Attention, Perception, & Psychophysics*, 71(7):1439–1459.

[Leithinger et al., 2011] Leithinger, D., Lakatos, D., DeVincenzi, A., Blackshaw, M., and Ishii, H. (2011). Direct and gestural interaction with relief: A 2.5D shape display. page 8.

[Lewis, 1993] Lewis, J. R. (1993). Multipoint scales: Mean and median differences and observed significance levels. *International Journal of Human-Computer Interaction*, 5(4):383–392.

[Li et al., 2010] Li, J., Martens, J.-B., and van Wijk, J. J. (2010). A model of symbol size discrimination in scatterplots. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2553–2562, Atlanta Georgia USA. ACM.

[Lie et al., 2009] Lie, A. E., Kehrer, J., and Hauser, H. (2009). Critical design and realization aspects of glyph-based 3D data visualization. In *Proceedings of the 25th Spring Conference on Computer Graphics*, pages 19–26, Budmerice Slovakia. ACM.

[López García and Hornecker, 2021] López García, I. and Hornecker, E. (2021). Scaling Data Physicalization – How Does Size Influence Experience? In *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction*, pages 1–14, Salzburg Austria. ACM.

[Losev et al., 2020] Losev, T., Storteboom, S., Carpendale, S., and Knudsen, S. (2020). Distributed synchronous visualization design: Challenges and strategies. In *2020 IEEE*

Workshop on Evaluation and Beyond-Methodological Approaches to Visualization (BELIV), pages 1–10. IEEE.

[Lozano-Hemmer, 2004] Lozano-Hemmer, R. (2004). Synaptic Caguamas: Visualize Cellular Automata. http://dataphys.org/list/synaptic-caguamas-visualize-cellular-automata/.

[Madsen, 2018] Madsen, L. (2018). New Worry Beads: Deaths from Terrorism. http://dataphys.org/list/new-worry-beads-global-deaths-from-terrorism/.

[McCormack et al., 2018] McCormack, J., Roberts, J. C., Bach, B., Freitas, C. D. S., Itoh, T., Hurter, C., and Marriott, K. (2018). Multisensory Immersive Analytics. In Marriott, K., Schreiber, F., Dwyer, T., Klein, K., Riche, N. H., Itoh, T., Stuerzlinger, W., and Thomas, B. H., editors, *Immersive Analytics*, volume 11190, pages 57–94. Springer International Publishing, Cham.

[Meier, 2017] Meier, S. (2017). Green Berlin. http://dataphys.org/list/green-berlin/.

[Merčun, 2014] Merčun, T. (2014). Evaluation of information visualization techniques: Analysing user experience with reaction cards. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*, pages 103–109, Paris France. ACM.

[Miebach, 2008] Miebach, N. (2008). Changing waters.

[Miebach et al., 2022] Miebach, N., Campbell, B. D., and Samsel, F. (2022). Nathalie Miebach: Sculpted Data Infused With Craftsmanship. *IEEE Computer Graphics and Applications*, 42(1):7–16.

[Milgram and Kishino, 1994] Milgram, P. and Kishino, F. (1994). A taxonomy of mixed reality visual displays. *IEICE TRANSACTIONS on Information and Systems*, 77(12):1321–1329.

[Millar et al., 2018] Millar, G. C., Tabrizian, P., Petrasova, A., Petras, V., Harmon, B., Mitasova, H., and Meentemeyer, R. K. (2018). Tangible Landscape: A Hands-on Method for Teaching Terrain Analysis. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–12, Montreal QC Canada. ACM.

[Mitasova et al., 2012] Mitasova, H., Harmon, R. S., Weaver, K. J., Lyons, N. J., and Overton, M. F. (2012). Scientific visualization of landscapes and landforms. *Geomorphology*, 137(1):122–137.

[Morris and Ebert, 2000] Morris, C. J. and Ebert, D. S. (2000). An Experimental Analysis of the Effectiveness of Features in Chernoff Faces.

[Munzner, 2008] Munzner, T. (2008). Process and Pitfalls in Writing Information Visualization Research Papers. In Kerren, A., Stasko, J. T., Fekete, J.-D., and North, C., editors, *Information Visualization*, volume 4950, pages 134–153. Springer Berlin Heidelberg, Berlin, Heidelberg.

[Munzner, 2009] Munzner, T. (2009). A Nested Model for Visualization Design and Validation. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):921–928.

[Munzner, 2014] Munzner, T. (2014). *Visualization Analysis and Design*. A K Peters/CRC Press, 0 edition.

[Munzner et al., 2006] Munzner, T., Johnson, C., Moorhead, R., Pfister, H., Rheingans, P., and Yoo, T. (2006). NIH-NSF visualization research challenges report summary. *IEEE Computer Graphics and Applications*, 26(2):20–24.

[Nagel and Pietsch, 2015] Nagel, T. and Pietsch, C. (2015). Cf. city flows: A comparative visualization of urban bike mobility.

[Nam et al., 2019] Nam, J. W., McCullough, K., Tveite, J., Espinosa, M. M., Perry, C. H., Wilson, B. T., and Keefe, D. F. (2019). Worlds-in-wedges: Combining worlds-in-miniature and portals to support comparative immersive visualization of forestry data. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 747–755. IEEE.

[Nielson et al., 1997] Nielson, G., Hagen, H., and Muller, H. (1997). Scientific visualization. Institute of Electrical & Electronics Engineers.

[Nilsson et al., 2020] Nilsson, N. C., Zenner, A., and Simeone, A. L. (2020). Haptic proxies for virtual reality: Success criteria and taxonomy. In *Workshop on Everyday Proxy Objects for Virtual Reality at CHI'20*.

[Nittala et al., 2015] Nittala, A. S., Li, N., Cartwright, S., Takashima, K., Sharlin, E., and Sousa, M. C. (2015). PLANWELL: Spatial user interface for collaborative petroleum well-planning. In *SIGGRAPH Asia 2015 Mobile Graphics and Interactive Applications*, SA '15, New York, NY, USA. Association for Computing Machinery.

[North and Shneiderman, 2000] North, C. and Shneiderman, B. (2000). Snap-together visualization: A user interface for coordinating visualizations via relational schemata.

In *Proceedings of the Working Conference on Advanced Visual Interfaces*, pages 128–135.

[Offenhuber, 2020] Offenhuber, D. (2020). Data by Proxy — Material Traces as Autographic Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):98–108.

[Omirou et al., 2016] Omirou, T., Perez, A. M., Subramanian, S., and Roudaut, A. (2016). Floating charts: Data plotting using free-floating acoustically levitated representations. In *2016 IEEE Symposium on 3D User Interfaces (3DUI)*, pages 187–190, Greenville, SC, USA. IEEE.

[Pahr, 2021] Pahr, D. (2021). *Vologram: Educational Craftworks for Volume Physicalization*. PhD thesis, Wien.

[Pasternak et al., 2017] Pasternak, E., Fenichel, R., and Marshall, A. N. (2017). Tips for creating a block language with blockly. In *2017 IEEE Blocks and beyond Workshop (B&B)*, pages 21–24. IEEE.

[Peck et al., 2013] Peck, J., Barger, V. A., and Webb, A. (2013). In search of a surrogate for touch: The effect of haptic imagery on perceived ownership. *Journal of Consumer Psychology*, 23(2):189–196.

[Peck and Childers, 2003] Peck, J. and Childers, T. L. (2003). Individual differences in haptic information processing: The "need for touch" scale. *Journal of consumer research*, 30(3):430–442.

[Petersen et al., 2019] Petersen, M. R., Asay-Davis, X. S., Berres, A. S., Chen, Q., Feige, N., Hoffman, M. J., Jacobsen, D. W., Jones, P. W., Maltrud, M. E., Price, S. F., et al. (2019). An evaluation of the ocean and sea ice climate of E3SM using MPAS and interannual CORE-II forcing. *Journal of Advances in Modeling Earth Systems*, 11(5):1438–1458.

[Petersen et al., 2015] Petersen, M. R., Jacobsen, D. W., Ringler, T. D., Hecht, M. W., and Maltrud, M. E. (2015). Evaluation of the arbitrary Lagrangian-Eulerian vertical coordinate method in the MPAS-Ocean model. *Ocean Modelling*, 86(0):93–113.

[Pezoa et al., 2016] Pezoa, F., Reutter, J. L., Suarez, F., Ugarte, M., and Vrgoč, D. (2016). Foundations of JSON schema. In *Proceedings of the 25th International Conference on World Wide Web*, pages 263–273.

[Pharr et al., 2016] Pharr, M., Jakob, W., and Humphreys, G. (2016). *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann.

[Pilar and Ware, 2013] Pilar, D. H. F. and Ware, C. (2013). Representing Flow Patterns by Using Streamlines with Glyphs. *IEEE Transactions on Visualization and Computer Graphics*, 19(8):1331–1341.

[Piper et al., 2002] Piper, B., Ratti, C., and Ishii, H. (2002). Illuminating clay: A 3-D tangible interface for landscape analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 355–362.

[Pousman et al., 2007] Pousman, Z., Stasko, J., and Mateas, M. (2007). Casual Information Visualization: Depictions of Data in Everyday Life. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1145–1152.

[Priestnall et al., 2012] Priestnall, G., Gardiner, J., Durrant, J., and Goulding, J. (2012). Projection Augmented Relief Models (PARM): Tangible Displays for Geographic Information. In *Electronic Visualisation and the Arts (EVA 2012)*.

[PTC, 2020] PTC (2020). Vuforia.

[Pustka et al., 2012] Pustka, D., Hülß, J.-P., Willneff, J., Pankratz, F., Huber, M., and Klinker, G. (2012). Optical outside-in tracking using unmodified mobile phones. In *2012 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 81–89. IEEE.

[Rase, 2011] Rase, W.-D. (2011). Creating Physical 3D Maps Using Rapid Prototyping Techniques. In Buchroithner, M., editor, *True-3D in Cartography*, pages 119–134. Springer Berlin Heidelberg, Berlin, Heidelberg.

[Ren and Hornecker, 2021] Ren, H. and Hornecker, E. (2021). Comparing Understanding and Memorization in Physicalization and VR Visualization. In *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction*, pages 1–7, Salzburg Austria. ACM.

[Resnick et al., 2009] Resnick, M., Maloney, J., Monroy-Hernández, A., Rusk, N., Eastmond, E., Brennan, K., Millner, A., Rosenbaum, E., Silver, J., Silverman, B., et al. (2009). Scratch: Programming for all. *Communications of the ACM*, 52(11):60–67.

[Rhyne, 2012] Rhyne, T.-M. (2012). Applying artistic color theories to visualization. *Expanding the Frontiers of Visual Analytics and Visualization*, pages 263–283.

[Riva et al., 2010] Riva, A., Conti, G., Solinas, P., and Loy, F. (2010). The evolution of anatomical illustration and wax modelling in Italy from the 16th to early 19th centuries. *Journal of Anatomy*, 216(2):209–222.

[Rocha et al., 2017] Rocha, A., Silva, J. D., Alim, U., and Sousa, M. C. (2017). Multivariate Visualization of Oceanography Data Using Decals. *Workshop on Visualisation in Environmental Sciences (EnvirVis)*, page 5 pages.

[Ropinski et al., 2011] Ropinski, T., Oeltze, S., and Preim, B. (2011). Survey of glyph-based visualization techniques for spatial multivariate medical data. *Computers & Graphics*, 35(2):392–401.

[Ropinski et al., 2007] Ropinski, T., Specht, M., Meyer-Spradow, J., Hinrichs, K., and Preim, B. (2007). Surface Glyphs for Visualizing Multimodal Volume Data.

[Saket et al., 2016] Saket, B., Endert, A., and Stasko, J. (2016). Beyond Usability and Performance: A Review of User Experience-focused Evaluations in Visualization. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization*, pages 133–142, Baltimore MD USA. ACM.

[Salisbury et al., 2004] Salisbury, K., Conti, F., and Barbagli, F. (2004). Haptic rendering: Introductory concepts. *IEEE Computer Graphics and Applications*, 24(2):24–32.

[Samsel et al.,] Samsel, F., Abram, G., Johnson, S., Herman, B., and Keefe, D. F. The Sculpting Vis Library.

[Samsel et al., 2018a] Samsel, F., Bartram, L., and Bares, A. (2018a). Art, Affect and Color: Creating Engaging Expressive Scientific Visualization. In *2018 IEEE VIS Arts Program (VISAP)*, pages 1–9, Berlin, Germany. IEEE.

[Samsel et al., 2019a] Samsel, F., Johnson, S., Bares, A., and Keefe, D. F. (2019a). Scientific Visualization: Enriching Vocabulary via the Human Hand.

[Samsel et al., 2018b] Samsel, F., Klaassen, S., and Rogers, D. H. (2018b). Colormoves: Real-time interactive colormap construction for scientific visualization. *IEEE computer graphics and applications*, 38(1):20–29.

[Samsel et al., 2015] Samsel, F., Petersen, M., Geld, T., Abram, G., Wendelberger, J., and Ahrens, J. (2015). Colormaps that Improve Perception of High-Resolution Ocean Data. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, pages 703–710, Seoul Republic of Korea. ACM.

[Samsel et al., 2019b] Samsel, F., Wolfram, P., Bares, A., Turton, T. L., and Bujack, R. (2019b). Colormapping resources and strategies for organized intuitive environmental visualization. *Environmental Earth Sciences*, 78(9):269.

[Saraiya et al., 2005] Saraiya, P., North, C., and Duca, K. (2005). An insight-based methodology for evaluating bioinformatics visualizations. *IEEE transactions on visualization and computer graphics*, 11(4):443–456.

[Satriadi et al., 2022] Satriadi, K. A., Smiley, J., Ens, B., Cordeil, M., Czauderna, T., Lee, B., Yang, Y., Dwyer, T., and Jenny, B. (2022). Tangible Globes for Data Visualisation in Augmented Reality. In *CHI Conference on Human Factors in Computing Systems*, pages 1–16, New Orleans LA USA. ACM.

[Satyanarayan et al., 2016] Satyanarayan, A., Moritz, D., Wongsuphasawat, K., and Heer, J. (2016). Vega-lite: A grammar of interactive graphics. *IEEE transactions on visualization and computer graphics*, 23(1):341–350.

[Sauvé et al., 2020] Sauvé, K., Potts, D., Alexander, J., and Houben, S. (2020). A change of perspective: How user orientation influences the perception of physicalizations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–12.

[Savage et al., 2012] Savage, V., Zhang, X., and Hartmann, B. (2012). Midas: Fabricating custom capacitive touch sensors to prototype interactive objects. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology*, pages 579–588.

[Schmandt-Besserat, 2013] Schmandt-Besserat, D. (2013). Tokens and writing: The cognitive development. In *UT Faculty/Researcher Works*. Ex Oriente.

[Schmitz, 2004] Schmitz, M. (2004). XenoVision Mark III: A Dynamic Solid Terrain Model. http://dataphys.org/list/xenovision-mark-iii-a-dynamic-solid-terrain-model/.

[Schroeder et al., 2010] Schroeder, D., Coffey, D. M., and Keefe, D. F. (2010). Drawing with the flow: A sketch-based interface for illustrative visualization of 2d vector fields. In *SBM*, pages 49–56. Citeseer.

[Schroeder and Keefe, 2016] Schroeder, D. and Keefe, D. F. (2016). Visualization-by-Sketching: An Artist's Interface for Creating Multivariate Time-Varying Data Visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):877–885.

[Schroeder et al., 2014] Schroeder, D., Korsakov, F., Knipe, C. M.-P., Thorson, L., Ellingson, A. M., Nuckley, D., Carlis, J., and Keefe, D. F. (2014). Trend-Centric Motion Visualization: Designing and Applying a new Strategy for Analyzing Scientific Motion Collections. *IEEE transactions on visualization and computer graphics*, 20(12):2644–2653.

[Schroeder et al., 2012] Schroeder, D., Kowalewski, T., White, L., Carlis, J., Santos, E., Sweet, R., Lendvay, T. S., Reihsen, T., and Keefe, D. (2012). Exploratory Visualization of Surgical Training Databases for Improving Skill Acquisition. *IEEE Computer Graphics and Applications*, 32(6):71–81.

[Schroeder et al., 2004] Schroeder, W. J., Lorensen, B., and Martin, K. (2004). *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*. Kitware.

[Schultz and Kindlmann, 2010] Schultz, T. and Kindlmann, G. L. (2010). Superquadric Glyphs for Symmetric Second-Order Tensors. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1595–1604.

[Schulze-Döbold et al., 2001] Schulze-Döbold, J., Wössner, U., Walz, S. P., and Lang, U. (2001). Volume rendering in a virtual environment. In *Immersive Projection Technology and Virtual Environments 2001*, pages 187–198. Springer.

[Schwanke, 1991] Schwanke, R. W. (1991). An intelligent tool for re-engineering software modularity. In *Proceedings-13th International Conference on Software Engineering*, pages 83–84. IEEE Computer Society.

[Segal, 2011] Segal, A. (2011). Adrien Segal's Data Furniture. http://dataphys.org/list/data-furniture/.

[Segal, 2015] Segal, A. (2015). Grewingk glacier.

[Sensel, Inc.,] Sensel, Inc. The Sensel Morph.

[Shaer and Hornecker, 2010] Shaer, O. and Hornecker, E. (2010). *Tangible User Interfaces: Past, Present, and Future Directions*. Now Publishers Inc.

[Shneiderman, 1996] Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings 1996 IEEE Symposium on Visual Languages*, pages 336–343.

[Shneiderman, 2000] Shneiderman, B. (2000). Creating creativity: User interfaces for supporting innovation. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(1):114–138.

[Shneiderman, 2002] Shneiderman, B. (2002). Creativity support tools. *Communications of the ACM*, 45(10):116–120.

[Shneiderman, 2004] Shneiderman, B. (2004). Designing for fun: How can we design user interfaces to be more fun? *interactions*, 11(5):48–50.

[Shneiderman, 2007] Shneiderman, B. (2007). Creativity support tools: Accelerating discovery and innovation. *Communications of the ACM*, 50(12):20–32.

[Sigitas Guzauskas, 2018] Sigitas Guzauskas (2018). Living Map: Precipitation Visualized with Moss. http://dataphys.org/list/living-map-precipitation-visualized-with-moss/.

[Silva et al., 2011] Silva, S., Sousa Santos, B., and Madeira, J. (2011). Using color in visualization: A survey. *Computers & Graphics*, 35(2):320–333.

[Slocum et al., 2022] Slocum, T. A., McMaster, R. B., Kessler, F. C., and Howard, H. H. (2022). *Thematic Cartography and Geovisualization*. CRC Press.

[Sobel et al., 2004] Sobel, J. S., Forsberg, A. S., Laidlaw, D. H., Zeleznik, R. C., Keefe, D. F., Pivkin, I., Karniadakis, G. E., Richardson, P., and Swartz, S. (2004). Particle flurries. *IEEE Computer Graphics and Applications*, 24(2):76–85.

[Song et al., 2011] Song, H., Benko, H., Guimbretiere, F., Izadi, S., Cao, X., and Hinckley, K. (2011). Grips and gestures on a multi-touch pen. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1323–1332.

[South et al., 2022] South, L., Saffo, D., Vitek, O., Dunne, C., and Borkin, M. A. (2022). Effective Use of Likert Scales in Visualization Evaluations: A Systematic Review. *Computer Graphics Forum*, 41(3):43–55.

[Spitz, 2013] Spitz, A. (2013). Loci: 3D Printed Sculptures of Your Flights. http://dataphys.org/list/loci-3d-printed-sculptures-of-your-flights/.

[Stasko, 2014] Stasko, J. (2014). Value-driven evaluation of visualizations. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*, pages 46–53, Paris France. ACM.

[Steed, 2008] Steed, A. (2008). A simple method for estimating the latency of interactive, real-time graphics simulations. In *Proceedings of the 2008 ACM Symposium on Virtual Reality Software and Technology*, pages 123–129.

[Steinicke et al., 2012] Steinicke, F., Benko, H., Krüger, A., Keefe, D., de la Riviére, J.-B., Anderson, K., Häkkilä, J., Arhippainen, L., and Pakanen, M. (2012). The 3rd dimension of CHI (3DCHI) touching and designing 3D user interfaces. In *CHI'12 Extended Abstracts on Human Factors in Computing Systems*, pages 2695–2698.

[Stevens et al., 2020] Stevens, A. H., Ware, C., Butkiewicz, T., Rogers, D., and Abram, G. (2020). Hairy Slices II: Depth Cues for Visualizing 3D Streamlines Through Cutting Planes. *Computer Graphics Forum*, 39(3):25–35.

[Stevenson et al., 2010] Stevenson, A., Perez, C., and Vertegaal, R. (2010). An inflatable hemispherical multi-touch display. In *Proceedings of the Fifth International Conference on Tangible, Embedded, and Embodied Interaction*, pages 289–292.

[Stoakley et al., 1995] Stoakley, R., Conway, M. J., and Pausch, R. (1995). Virtual reality on a WIM: Interactive worlds in miniature. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 265–272.

[Stusak et al., 2014] Stusak, S., Tabard, A., Sauka, F., Khot, R. A., and Butz, A. (2014). Activity Sculptures: Exploring the Impact of Physical Visualizations on Running Activity. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2201–2210.

[Swackhamer et al., 2017] Swackhamer, M., Johnson, A. J., Keefe, D., Johnson, S., Altheimer, R., and Wittkamper, A. (2017). Weather report: Structuring data experience in the built environment. *Proceedings of Architectural Research Centers Consortium*, pages 102–111.

[Taher et al., 2015] Taher, F., Hardy, J., Karnik, A., Weichel, C., Jansen, Y., Hornbæk, K., and Alexander, J. (2015). Exploring Interactions with Physically Dynamic Bar Charts. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 3237–3246, Seoul Republic of Korea. ACM.

[Taher et al., 2016] Taher, F., Jansen, Y., Woodruff, J., Hardy, J., Hornbaek, K., and Alexander, J. (2016). Investigating the Use of a Dynamic Physical Bar Chart for Data Exploration and Presentation. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):451–460.

[Tang et al.,] Tang, S. K., Sekikawa, Y., Leithinger, D., Follmer, S., and Ishii, H. Tangible CityScape.

[Tateosian et al., 2010] Tateosian, L., Mitasova, H., Harmon, B., Fogleman, B., Weaver, K., and Harmon, R. (2010). TanGeoMS: Tangible Geospatial Modeling System. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1605–1612.

[Taylor et al., 2001] Taylor, R. M., Hudson, T. C., Seeger, A., Weber, H., Juliano, J., and Helser, A. T. (2001). VRPN: A device-independent, network-transparent VR peripheral system. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, pages 55–61.

[Tejada et al., 2020] Tejada, C. E., Ramakers, R., Boring, S., and Ashbrook, D. (2020). AirTouch: 3D-printed touch-sensitive objects using pneumatic sensing. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–10.

[Terry and Mynatt, 2002] Terry, M. and Mynatt, E. D. (2002). Side views: Persistent, on-demand previews for open-ended tasks. In *Proceedings of the 15th Annual ACM Symposium on User Interface Software and Technology*, pages 71–80.

[Thrun et al., 2016] Thrun, M. C., Lerch, F., Lötsch, J., and Ultsch, A. (2016). Visualization and 3D Printing of Multivariate Data of Biomarkers. page 10.

[Tong et al., 2022] Tong, W., Chen, Z., Xia, M., Lo, L. Y.-H., Yuan, L., Bach, B., and Qu, H. (2022). Exploring Interactions with Printed Data Visualizations in Augmented Reality.

[Tory et al., 2006] Tory, M., Kirkpatrick, A., Atkins, M., and Moller, T. (2006). Visualization task performance with 2D, 3D, and combination displays. *IEEE Transactions on Visualization and Computer Graphics*, 12(1):2–13.

[Ulbricht et al., 2006] Ulbricht, C., Wilkie, A., and Purgathofer, W. (2006). Verification of Physically Based Rendering Algorithms. *Computer Graphics Forum*, 25(2):237–255.

[Upson et al., 1989] Upson, C., Faulhaber, T., Kamins, D., Laidlaw, D., Schlegel, D., Vroom, J., Gurwitz, R., and Van Dam, A. (1989). The application visualization system: A computational environment for scientific visualization. *IEEE Computer Graphics and Applications*, 9(4):30–42.

[Valcasara, 2015] Valcasara, N. (2015). *Unreal Engine Game Development Blueprints*. Packt Publishing Ltd.

[van Dam et al., 2002] van Dam, A., Laidlaw, D. H., and Simpson, R. M. (2002). Experiments in Immersive Virtual Reality for Scientific Visualization. *Computers & Graphics*, 26(4):535–555.

[van Onzenoodt et al., 2022] van Onzenoodt, C., Vázquez, P.-P., and Ropinski, T. (2022). Out of the Plane: Flower Vs. Star Glyphs to Support High-Dimensional Exploration in Two-Dimensional Embeddings. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–15.

[van Wijk, 2005] van Wijk, J. J. (2005). The Value of Visualization. pages 79–86.

[Venkatesan et al., 2021] Venkatesan, M., Mohan, H., Ryan, J. R., Schürch, C. M., Nolan, G. P., Frakes, D. H., and Coskun, A. F. (2021). Virtual and augmented reality for biomedical applications. *Cell Reports Medicine*, 2(7):100348.

[Vepakomma, 2014] Vepakomma, M. (2014). *Blender Compositing and Post Processing*. Packt Publishing Ltd.

[Viégas and Wattenberg, 2007] Viégas, F. B. and Wattenberg, M. (2007). Artistic data visualization: Beyond visual analytics. In *International Conference on Online Communities and Social Computing*, pages 182–191. Springer.

[Walker and Nowacki, 2011] Walker, E. and Nowacki, A. S. (2011). Understanding Equivalence and Noninferiority Testing. *Journal of General Internal Medicine*, 26(2):192–196.

[Wang et al., 2008] Wang, L., Giesen, J., McDonnell, K. T., Zolliker, P., and Mueller, K. (2008). Color Design for Illustrative Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1739–1754.

[Wang et al., 2019] Wang, Y., Segal, A., Klatzky, R., Keefe, D. F., Isenberg, P., Hurtienne, J., Hornecker, E., Dwyer, T., and Barrass, S. (2019). An Emotional Response to the Value of Visualization. *IEEE Computer Graphics and Applications*, 39(5):8–17.

[Ware, 1988] Ware, C. (1988). Color sequences for univariate maps: Theory, experiments and principles. *IEEE Computer Graphics and Applications*, 8(5):41–49.

[Ware et al., 1993] Ware, C., Arthur, K., and Booth, K. S. (1993). Fish tank virtual reality. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '93*, pages 37–42, Amsterdam, The Netherlands. ACM Press.

[Weiß et al., 2020] Weiß, M., Angerbauer, K., Voit, A., Schwarzl, M., Sedlmair, M., and Mayer, S. (2020). Revisited: Comparison of empirical methods to evaluate visualizations supporting crafting and assembly purposes. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1204–1213.

[Willis, 2018] Willis, S. (2018). The Maker Revolution. *Computer*, 51(3):62–65.

[Wilson, 2010] Wilson, A. D. (2010). Using a depth camera as a touch sensor. In *ACM International Conference on Interactive Tabletops and Surfaces*, pages 69–72.

[Wilson, 2002] Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9(4):625–636.

[Wixon, 2011] Wixon, D. (2011). Measuring fun, trust, confidence, and other ethereal constructs: It isn't that hard. *Interactions*, 18(6):74–77.

[Wolfram et al., 2015] Wolfram, P. J., Ringler, T. D., Maltrud, M. E., Jacobsen, D. W., and Petersen, M. R. (2015). Diagnosing isopycnal diffusivity in an eddying, idealized midlatitude ocean basin via Lagrangian, in Situ, Global, High-Performance Particle Tracking (LIGHT). *Journal of Physical Oceanography*, 45(8):2114–2133.

[Woods et al., 2016] Woods, T. L., Reed, S., Hsi, S., Woods, J. A., and Woods, M. R. (2016). Pilot study using the augmented reality sandbox to teach topographic maps and surficial processes in introductory geology labs. *Journal of Geoscience Education*, 64(3):199–214.

[Yi et al., 2008] Yi, J. S., Kang, Y.-a., Stasko, J. T., and Jacko, J. A. (2008). Understanding and characterizing insights: How do people gain insights using information visualization? In *Proceedings of the 2008 Workshop on BEyond Time and Errors: Novel evaLuation Methods for Information Visualization*, pages 1–6.

[Zeleznik et al., 2002] Zeleznik, R. C., LaViola, J. J., Feliz, D. A., and Keefe, D. F. (2002). Pop through button devices for VE navigation and interaction. In *Proceedings IEEE Virtual Reality 2002*, pages 127–134. IEEE.

[Zeller et al., 2022] Zeller, S., Samsel, F., and Bartram, L. (2022). Affective, Hand-Sculpted Glyph Forms for Engaging and Expressive Scientific Visualization. In *2022 IEEE VIS Arts Program (VISAP)*, pages 127–136.

[Zhang et al., 2016] Zhang, C., Schultz, T., Lawonn, K., Eisemann, E., and Vilanova, A. (2016). Glyph-Based Comparative Visualization for Diffusion Tensor Fields. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):797–806.

[Zhang et al., 2020] Zhang, G., Gong, J., Li, Y., Sun, J., Xu, B., Zhang, D., Zhou, J., Guo, L., Shen, S., and Yin, B. (2020). An efficient flood dynamic visualization approach based on 3D printing and augmented reality. *International Journal of Digital Earth*, 13(11):1302–1320.

[Zhou and Hansen, 2016] Zhou, L. and Hansen, C. D. (2016). A Survey of Colormaps in Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 22(8):2051–2069.

[Zhu, 2013] Zhu, N. Q. (2013). *Data Visualization with D3. Js Cookbook*. Packt Publishing Ltd.

[Zimmerman and Forlizzi, 2014] Zimmerman, J. and Forlizzi, J. (2014). Research Through Design in HCI. In Olson, J. S. and Kellogg, W. A., editors, *Ways of Knowing in HCI*, pages 167–189. Springer New York, New York, NY.