TU WIEN Informatics

From Interactions to Integrated Actions

Exploring Active Perception and Inter-Modality in Data Physicalization

DISSERTATION

zur Erlangung des akademischen Grades

Doktor der Technischen Wissenschaften

eingereicht von

DI Daniel Pahr, Bsc
Matrikelnummer 0906438

an der Fakultät für Informatik

der Technischen Universität Wien

Betreuung: Associate Prof. Dr. Renata G. Raidou
Zweitbetreuung: Hsiang-Yun Wu, PhD

Diese Dissertation haben begutachtet:

Helen Purchase

Henry Fuchs

Wien, August 12, 2025

Daniel Pahr

TU WIEN Informatics

From Interactions to Integrated Actions

Exploring Active Perception and Inter-Modality in Data Physicalization

DISSERTATION

submitted in partial fulfillment of the requirements for the degree of

Doktor der Technischen Wissenschaften

by

DI Daniel Pahr, Bsc
Registration Number 0906438

to the Faculty of Informatics

at the TU Wien

Advisor: Associate Prof. Dr. Renata G. Raidou
Second advisor: Hsiang-Yun Wu, PhD

The dissertation has been reviewed by:

Helen Purchase	Henry Fuchs

Vienna, August 12, 2025

Daniel Pahr

Technische Universität Wien
A-1040 Wien ▪ Karlsplatz 13 ▪ Tel. +43-1-58801-0 ▪ www.tuwien.at

Erklärung zur Verfassung der Arbeit

DI Daniel Pahr, Bsc

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.
Ich erkläre weiters, dass ich mich generativer KI-Tools lediglich als Hilfsmittel bedient habe und in der vorliegenden Arbeit mein gestalterischer Einfluss überwiegt. Im Anhang „Übersicht verwendeter Hilfsmittel“ habe ich alle generativen KI-Tools gelistet, die verwendet wurden, und angegeben, wo und wie sie verwendet wurden. Für Textpassagen, die ohne substantielle Änderungen übernommen wurden, haben ich jeweils die von mir formulierten Eingaben (Prompts) und die verwendete IT- Anwendung mit ihrem Produktnamen und Versionsnummer/Datum angegeben.

Wien, August 12, 2025

Daniel Pahr

Acknowledgements

I don't think five years ago I would have ever even seen myself ever writing a dissertation, but since then, my life and career have substantially changed.

This new path for my life was made possible by Renata and Yun, who gave me the chance to start over and supported me until the very end of this crazy endeavor. The past four years, the vis group at TU Wien became a family to me, and the confusing and smelly halls of Favoritenstrasse 9 became a home. Even as I depart for new endeavors, I will always remember the much-too-long coffee breaks with Joey, Henry, Nidham, and all the others I frequently kept from work, and bring this spirit of camaraderie to wherever I go.

I also learned that our work is not done by one person and that I will always have my trusty co-authors at my side. Without Manu, Michal, Velitchko, and all the other names that are now a matter of publication anyway, none of this work would have seen the light of day. And to Sara, who came into my life at the right time and in the weirdest possible place, and who never stopped believing in me, I owe the greatest debt. My parents and my brother have my eternal gratitude for always supporting me, especially when I needed someone to take care of my cat.

Finally, my trusty companion, Galvina, the angry creature, who has always forgiven me abandoning her temporarily, will always have a place on my lap.

Kurzfassung

Das wachsende Feld der Datenphysikalisierung birgt signifikantes Potential für die Integration des Handelns der Nutzer durch physikalische Artefakte. Zwei vielversprechende Faktoren für solche physischen, im Gegensatz zu virtuellen Repräsentationen sind die *physikalischen Interaktionen* und die *multimodale Wahrnehmung*. Direkte Manipulation im physikalischen Raum erlaubt Nutzern, Datenphysikalisierungen auf natürliche Art zu erforschen, und das Handeln des Nutzers zum Kodieren und Dekodieren von Informationen zu verwenden, anders als bei rein virtuellen Repräsentationen. In dieser Dissertation erforsche ich das Einbinden der Handlungen von Nutzern, sowohl im bezug auf Interaktion als auch Wahrnehmung, wobei ich mich zwischen Repräsentationen wo Interaktionen rein vor der Wahrnehmung stattfinden und solchen wo physikalische Interaktionen die menschlichen Sinne direkt stimulieren.

Ich untersuche vier verschiedene Arten der Interaktion mit Datenphysikalisierungen und zeige, wie jede davon die menschlichen Wahrnehmungen unterschiedlich beeinflusst. Erstens zeige ich, wie eine **modulare** 3D Darstellung von dynamischen Daten durch physikalische Verkörperung die natürliche Wahrnehmung von Menschen nutzen kann. Dafür präsentiere ich eine einfache interaktive physikalische Repräsentation eines Raumzeit Würfels, die ich in einer Fallstudie mit einer Domänenexpertin bespreche. Zweitens untersuche ich den Einfluss des Aktes der **Konstruktion** — eine intuitive Interaktion im physikalischen Raum — einer vorgegebenen Repräsentation auf die menschliche Wahrnehmung. Ich zeige dies anhand des Designs eines Netzwerk-Physikalisierungsbaukastens und dessen Einsatz in einer Studie, wo ich verschiedene Methoden, eine Visualisierung zu konstruieren, vergleiche. Drittens führe ich die **taktile** Wahrnehmung von elastischen Eigenschaften eines Objekts durch eine multimodale Repräsentation von Volumendaten ein. Dafür stelle ich eine Fabrikations-Pipeline vor, mit der Nutzer elastische Artefakte aus Volumendaten mit konventionellem 3D Druck erzeugen können und validiere diese durch verschiedene Methoden. Letztlich untersuche ich den Nutzen der manuellen **Bedienung** eines dynamischen Prozesses, die die Wahrnehmung eines Nutzers direkt stimuliert, um Informationen zu liefern. In einer Studie zeige ich, dass die positiven Effekte der Integration des Handelns eines Nutzers in eine Repräsentation.

Insgesamt zeigen die Resultate, dass bereits eine einfache Datenphysikalisierung durch natürliche menschliche Wahrnehmung positive Effekte erzielen kann. Abstrakte Repräsentationen müssen von den Nutzern erlernt werden, während physikalische Metaphern durch direkte Interaktion von der menschlichen Wahrnehmung profitieren, sofern der Reiz auf die Fähigkeiten der Sinne angepasst ist.

Abstract

The growing field of data physicalization holds significant potential for integrating **user actions** directly into the sensemaking process through physical artifacts. Two promising factors for physical, as opposed to virtual representations, are *physical interaction* and *multimodal perception*. Unmediated interaction in the physical space allows users to manipulate and explore data physicalizations in a natural way, harnessing a user's actions to encode and decode information in a different way than purely virtual representations. In this dissertation, I explore the incorporation of user action as a means of manipulation and perception into data physicalizations, moving from representations where perception only happens after physical interactions, to representations where physical interactions directly stimulate the user's perception.

I investigate four distinct types of user interactions with data physicalizations and show how each of them can support human perception in different ways. Firstly, I show how a **modular** 3D representation of dynamic data can leverage physical embodiment using natural spatial perception. I demonstrate this by creating a simple interactive physical representation of a space-time-cube metaphor and investigating it in a case study with a domain expert. Secondly, I investigate the influence of **construction** — an intuitively physical interaction in the physical space — of a pre-defined physical representation on human perception. I show this by designing a network data physicalization toolkit and conducting a between-subject study, comparing different ways to instruct a user during construction. Thirdly, I introduce **tactile** perception of the elastic properties of an object in a multi-modal representation of volume data. I showcase this at the hands of a fabrication pipeline that creates elastic artifacts from volume data using consumer-level 3D printing and validate the method through computational, mechanical, and perceptual studies. Finally, I explore the benefits of manually **operating** a physical representation of a dynamic process, leveraging the tactile feedback to the user for information encoding. By means of a between-subject user study, I show that integrating a user's actions into a representation significantly increases engagement.

Overall, the results show that even a simple physicalization can highlight the perceptual benefits of physically encoding data by ways of natural perception. Abstract representations have to be learned by users but can be supported by physical interactions, while embodied metaphors profit from direct interactivity if the stimulus fits the sensory capabilities.

Contents

CHAPTER 1

Introduction

Ever since the early days of humanity, man-made artifacts have helped us make sense of the world around us. Historically, we find clay tokens used for accounting and trade (Figure 1.1a), abacuses for calculations, and sculptures representing human anatomy in many cultures. Before the digital age, scientific work was still conducted on blackboards and paper. Three-dimensional data representations, such as Watson and Crick's famous DNA model (Figure 1.1b), were conceived in physical space. There, they supported researchers' reasoning about phenomena, while teachers explained three-dimensional concepts to their students [DJVM20].

With the onset of computers and the ability to manipulate graphical displays on the fly, screen-based visualizations have become indispensable and have replaced their traditional physical counterparts. Rapidly, data visualizations grew from static charts into interactive knowledge-enhancing tools [CMS99]. The relative ease with which we could suddenly create visual representations and how we can alter, replace, or refine them interactively has given birth to the field of *data visualization*. Today, we find visualizations as inherent components of analytics and knowledge communication in many forms — from desktop applications to immersive head-mounted displays. As the name implies, data visualization, by default, relies on visual perception alone. This may be owed to the fact that, on the one hand, human vision is very precise and accurate, and, on the other hand, (pixel-based) computer screens are a convenient method to display even continuously updating and growing information.

Human perception, however, comprises the entire spectrum of senses. The theory of embodied cognition states that human perception is shaped by interactions with the outside world [NEFM99]. Besides visual representations, data could be turned into sounds, tactile sensations, tastes, or even smells. Instead of focusing purely on the visual component of perception, the term *data physicalization* [JDI+15] encompasses this notion, describing how data representations such as bar charts (Figure 1.1c) can be brought into the real world. In the last decade, bringing data into the physical space has shown potential in ways that are neglected or simply out of scope for data visualization, such as unique real-life interactions and specific ways to address human perception. We can naturally investigate physical artifacts, both in terms of *interactivity* (e.g.,

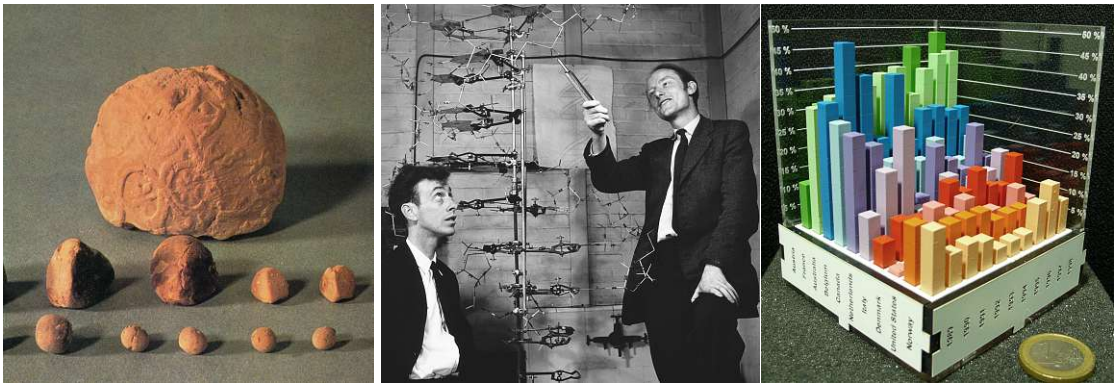

(a) (b) (c)

Figure 1.1: Examples of physical data representations throughout history. (a) Clay tokens from 5500 BC used as externalization tools for counting. (b) Watson and Crick examine their physical model of the DNA helix. (c) A digitally fabricated 3D-bar chart by Jansen et al. [JDF13]. Images taken from `dataphys.org` [DJ12].

picking something up and rotating it) and *perception* (e.g., using stereoscopic vision to examine a real-life artifact) [JDI+15].

While screens have become very precise, haptic feedback devices are not a common addition to a desktop setup. Interactions with visualizations are well supported through mouse, keyboard, and even touch screens. Compared to virtual representations, in a physical representation, there is no need for mediation of feedback over screens, haptic devices, or headphones. In this dissertation, I highlight the benefits of physical data representations with a focus on active perception and multimodality. I propose that data physicalizations provide a wide spectrum of possibilities for leveraging user actions for information encoding. My research shows this through a range of examples, from **modular physicalizations**, where users interactively rearrange components to guide visual inspection, to **manually operated representations**, where the observer's actions are inseparable from the functioning of the model itself. These cases illustrate a shift from **interaction** to **integrated action**. Figure 1.2 shows a sketch of this spectrum at the hand of various examples.

1.1 Definitions and Problem Statement

I investigate the design space of data physicalization by focusing on two factors, *interactivity* and *multi-modality*. Here, I describe how these factors contribute to the integration of user action into non-physical and physical data representations.

In their user-action framework for human-computer interaction, Hartson [Har03] proposes a four-stage cycle for user interaction. At the start of the cycle, the user **assesses** the outcomes of the previous iteration, i.e., the state of the physicalization. This leads to the user planning the next action and translating it into the specific "language" of the interface, for example, clicking a

button to sort data, or deforming a physical artifact. Then, the actual **physical action** follows, which constitutes a manipulation of the representation. In the case of a virtual visualization rendering the result of the action, this leads back into the assessment stage. Bae et al. [BZW+22] define this as **indirect interactivity**. If the action causes a sensory action before the next cycle iteration, like in the case of the user deforming an artifact, where the user would feel the resistance of the internal structure, this constitutes **direct interactivity**.

Data visualization offers great potential for *indirect interactivity*. Depending on the purpose of visualization, users can **manipulate** the representation through, for example, selecting interesting data points or filtering and aggregation of values [BM13]. Changing the representation, like switching types of charts on the fly, can be easily done if the outcome is a pixel-based rendering. Schneiderman's [Shn03] Mantra, *"Overview first, zoom and filter, then details-on-demand."* is deeply ingrained in visualization design and well supported by modern digital visualization platforms. Providing such interactivity through a physical artifact is more complicated. Digital manufacturing methods, such as fused filament fabrication (FFF), the most popular form of **3D printing**, are intended to create solid artifacts. The behavior of solid matter is much harder to influence than pixels on a screen; as such, even filtering a dataset becomes difficult after the artifact is created. Technical advancements, like **fabrication techniques** and actuation, are the first step in the direction of enabling interactivity in physicalizations, for example, 3D printing with elastic materials or using a motor to animate a sculpture.

Physical artifacts also introduce opportunities for direct interactivity. When decoding the information in a physicalization requires user action, this can be understood as active perception. Here, I define **perception** as decoding the information encoded in a physicalization, and **sensing** as engaging human senses during the act of perception. While multiple senses can be used for encoding, my focus is on visual and tactile modalities. In this dissertation, **tactile** refers to unmediated interaction (e.g., touching an object), while **haptic** refers to simulated feedback (e.g., vibrations generated by a device to indicate physical resistance).

Card et al. [CMS99] define visualization as "*the use of computer-supported interactive, **visual** representations of data.*" to amplify *perception*. They describe visualization as a supporting tool for human cognition, leveraging the relative ease and accessibility of using computer graphics. To this day, the screen is the most used and most versatile output device on any computer. Other ways of information encoding have developed alongside visual representation, for example, **sonification** [EEC+24], which employs sound as the channel to encode information. **Tactile encodings** are often connected to accessibility, calling to mind Braille writing. Data physicalizations inherently possess the potential to engage a user on multiple sensory channels.

Multi-modal approaches combine multiple sensory channels, effectively broadcasting information. Haptic feedback in virtual reality (VR) simulation, for example, can communicate mechanical resistance to the user directly, in addition to the visual representation. When instead multiple sensory channels are necessary to encode a single attribute, we speak of **multisensory displays** [HH16]. Jansen et al. [JDI+15] refer to this as **inter-modal** perception.

Figure 1.2 shows a variety of interactions afforded by physical artifacts and how these can aid users in decoding their meaning. A modular physicalization, where the parts of the final representation

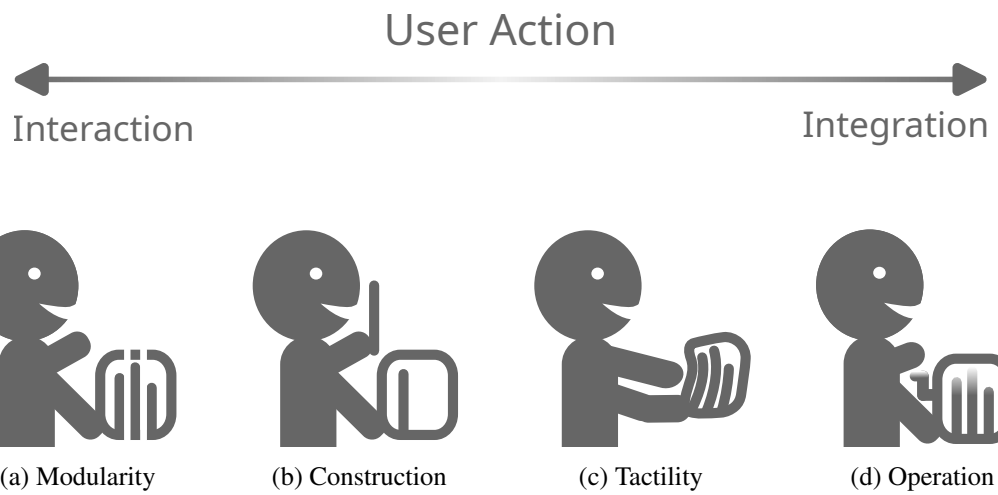

(a) Modularity (b) Construction (c) Tactility (d) Operation

Figure 1.2: Different examples of physical actions being afforded by physicalizations. (a) A modular physicalization allows the user to rearrange parts, and the interactivity facilitates visual inspection. (b) The user constructs a physicalization, while tangible interactions shape the representation, the information is only perceived visually. (c) The user explores an artifact with their hands, while information is perceived in parallel on the visual channel. (d) The user operates an active physicalization, integrating their actions into the representation.

are created in advance (Figure 1.2a), affords interactions similar to non-physical visualizations, such as filtering a dataset or isolating data points. The user manipulates the representation before decoding the data, only using their visual sense. When constructing a physical artifact, for example, by assembling a diagram using pre-defined building blocks (Figure1.2b), tactile interactions take place even before the representation is completed. Still, the data is only encoded visually. We can also encode data multi-modally, for example, by using haptic feedback through additional output devices on computers. Humans are accustomed to examining things both tangibly and visually (Figure 1.2c), for example, when investigating the surface of an object or pressing a fruit to determine its ripeness. The tactile interaction and sensing take place at the same time, and the user can visually inspect the object in addition. However, the stimuli represent the same data redundantly, making such a representation multi-modal. Information can also be encoded in the dynamic behavior of an object, for example, the process of transporting water through a mechanical pump. The user operates a pump, and senses the result of their action tactually in the mechanical resistance and visually in the flow of liquid transported through the apparatus. The resistance of the pump handle does not correlate with the flow of liquid; the complete picture can only be seen with both stimuli in parallel, making the representation inter-modal. The representation is not complete without the user operating it; the user's actions are integrated in the artifact, making the user an integral part of the representation (Figure 1.2d).

Screens have already evolved vastly since their inception, fabrication becomes more affordable every day, and VR devices have found their way into many homes. With the ongoing evolution of ever more immersive technologies in the digital sector, it becomes even more important to explore

how human interactions can be effectively utilized. There are benefits and pitfalls in different types of user interactions with physical data representations. My research during my time in the Ph.D. program at TU Wien revolved around exploring this design space, for which I present four distinct examples — modular representations (Figure 1.2a), construction (Figure 1.2b), tactile perception (Figure 1.2c), and manual operation (Figure 1.2d) — to study the different ways data physicalization can support sense-making.

> In this dissertation, I propose that using active perception and inter-modal encodings facilitates the integration of human actions into the sensemaking process.

1.2 Objective and Contributions

The objective of this dissertation is to highlight the interplay between active perception and inter-modality in the context of data physicalization and how this creates unique opportunities to include a user's actions in the sense-making process — **from interaction to integration** (Figure 1.2). To achieve this, I pose four distinct research questions, aimed at different ways a user can interact with a data physicalization.

Q1: ***Which benefits arise from physically representing a 3D visualization of dynamic data with a modular representation?*** Data visualization makes use of abstract concepts to help users analyze dynamic data. These concepts sometimes rely on three-dimensional metaphors, for example, a space-time cube visualization of time-dependent data. On screens, these metaphors can be communicated through 3D rendering. However, it has been shown that three-dimensional visualizations can suffer from a multitude of issues resulting from the two-dimensional projection needed to render them on a flat surface, such as occlusion and distortion effects. If instead rendered in physical space, where we can use *natural* depth perception on a three-dimensional object (Figure 1.2a), these shortcomings can be mitigated. Although this has been shown in the context of bar charts in the past [JDF13], it was also shown that interactive 2D visualization is more effective. I investigate a space-time cube, where the third dimension is used to encode the time axis, and individual parts can be rearranged in a modular representation.

Q2: ***What is the influence of constructing a physicalization with a pre-defined visual embedding on human perception?*** Knowledge externalization and its benefit to sensemaking is a central concept in data visualization [CMS99]. Interactions with virtual systems are highly sophisticated and provide a clear benefit to analysis tasks. Such interactions also benefit from the fact that computer graphics can be quickly generated and updated. Fabrication of physical representations, on the other hand, is a different, often more time-consuming process. However, this also comes with an opportunity. Embedded cognition suggests that human perception is grounded in the real world; thus, physical interactions with data have the potential to facilitate unique sensemaking processes [NEFM99] (Figure 1.2b). Specifically, for network diagrams, the act of creating a layout is an interesting aspect because the layout significantly influences the readability of the

final representation. I investigate the process of network layout creation in the physical space and compare how users interact with freely laid-out graphs vs. predetermined layouts.

Q3: ***How does tactile perception of local elastic compliance enhance the exploration of three-dimensional artifacts?*** Physicalizations offer a more natural way to interact with a representation than traditional input and output devices, such as mice, keyboards, and monitors. In the real world, we use our hands to naturally interact with things to investigate them. People naturally pick things up and use their fingers to help navigate structures. Arguably, there is a clear difference between physically interacting with an object and using our tactile senses to perceive it. Often, surface properties like texture are used to encode additional data into physical artifacts, for example, using roughness to indicate additional data on relief maps [DMSJ23]. While in real life, we often employ our tactile sense in addition to visual inspection, like inspecting an avocado in the supermarket or doctors using palpation on patients, little is known about the potential of the ability of sub-surface properties to encode data (Figure 1.2c). I investigate this by proposing a mapping of scalar field data to local elastic properties and conducting an investigation of how efficiently humans perceive the artifacts created with it.

Q4: ***Which are the benefits of manually operating a physical representation of a dynamic process compared to virtual and operated ones?*** Exploring the tactile properties of an artifact opens up an additional channel for information encoding. Yet, often, this will still encode information that can be perceived in visual encodings, such as the absorption of X-rays in computed tomography (CT) scans. The human body is capable of perceiving the environment using several senses at a time. This also enables us to reason about dynamic processes in abstract metaphors, such as visualizing energy consumption by making a user produce the same amount. While this requires us to make a step towards abstracting a complex process in such a way that a human can be involved, this also helps to put highly complex processes into understandable contexts. Using a representation where external influences have to be compensated can serve as an abstract representation. Learning theory has explored this, for example, in explaining mechanical physics with real-life apparatuses [FH95, GCC+10] (Figure 1.2d). The human body provides an interesting example with many dynamic processes that can be affected by external conditions, such as the cardiac cycle and the influence of vascular disease. To address this question, I design an abstract physical metaphor for the operation of the human heart and investigate how it enables users to perceive different abnormal states.

> The contribution of this dissertation comprises the design and execution of methodologies and accompanying studies that investigate the benefits of different ways to integrate user action into data physicalization.

The projects and their corresponding contributions are presented below, starting from purely visual and indirectly interactive representation and progressively adding tactile interactions, moving to direct interactivity.

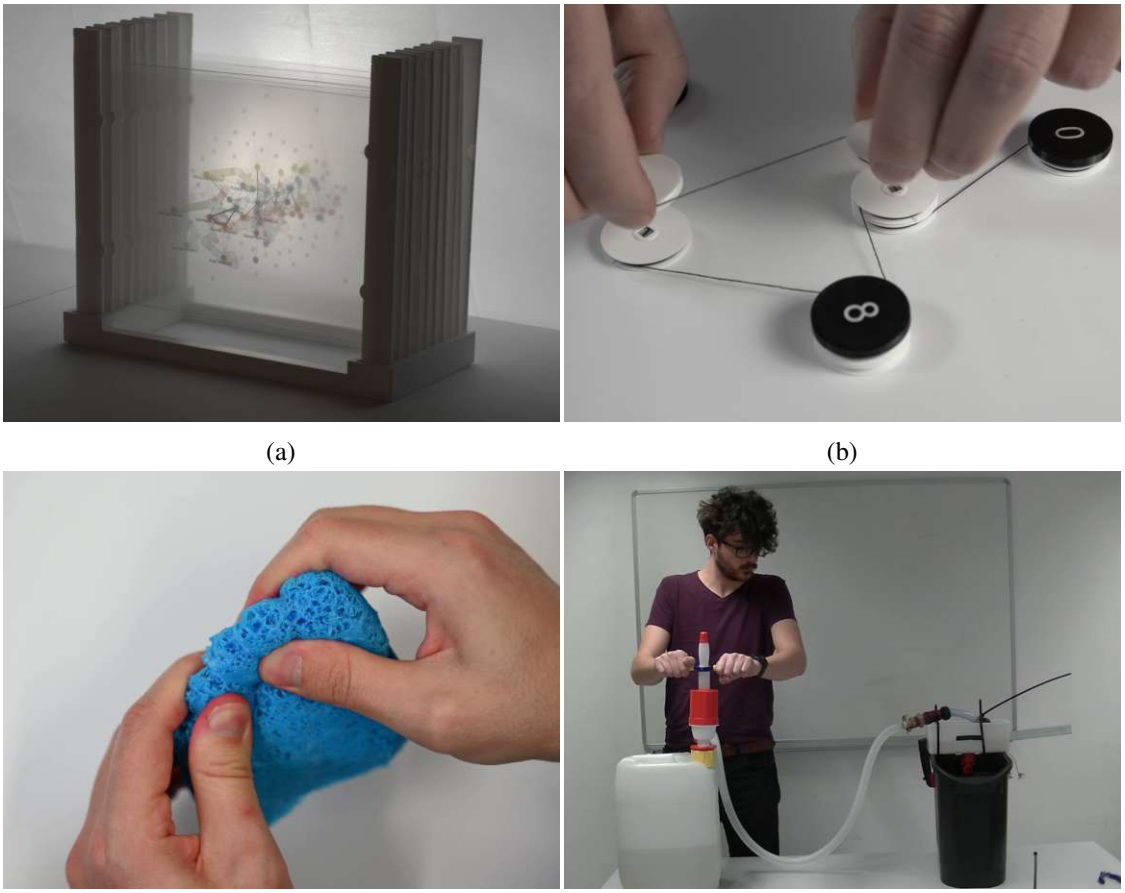

(a) (b)

(c) (d)

Figure 1.3: Physicalizations investigated in this dissertation with different physical interaction capabilities. (a) A modular 3D visualization that only communicates using the visual channel. (b) Constructing a physicalization adds tactile interaction, without tactile sensing. (c) Tactile exploration adds an additional channel to purely visual inspection. (d) Operating an interactive representation involves multiple senses in parallel to perceive the information.

P_1**: HoloGraphs** In Chapter 3, we present an interactive data physicalization for dynamic networks (Figure 1.3a). The purpose of *HoloGraphs* is to provide users with an affordable and accessible tool to create a modular space-time-cube representations of time-sliced dynamic graphs. Furthermore, we provide interaction in the form of overlays that can be added and removed by users. This constitutes a tactile interaction, but users observe data representation purely visually. ***Contribution:*** *We show that users' understanding of node trajectories around forming and dissolving group structures in social networks can profit from "natural" 3D perception.*

P_2**: NODKANT** Constructing a personalized representation of a dataset can lead to sense-making processes occurring even before it is completed. The embedding of a network diagram

can significantly contribute to its readability. Usually, these embeddings are computed by an algorithm, and users are only presented with the final network diagram. In Chapter 4, we design a versatile network physicalization toolkit and study the long and short-term effects on the understanding of network structures in a lab study (Figure 1.3b). ***Contribution:*** *We show that the construction of the embedding of a node-link diagram provides users with deeper and more memorable insights than interaction with a pre-constructed representation.*

P_3: Squishicalization With *squishicalization*, we present an interactive data physicalization pipeline to create a tangible representation of volume data. Here, we go beyond surface properties and use internal structures and elastic materials to control local elasticity, based on the scalar distribution (Figure 1.3c). In Chapter 5, we evaluate the technical feasibility, as well as the usability of our approach in a lab study. ***Contribution:*** *We show that the encoding of scalar field data into the elastic properties of an object stimulates tactile perception and provides insights beyond purely visual inspection.*

P_4: Heart Machine Finally, we present a data physicalization that embodies the cardiac cycle as an abstract metaphor, represented by a pump. In Chapter 6, we create four different versions of this representation in a two-dimensional design space describing the mode of operation and manifestation (Figure 1.3d). In a mixed-methods lab study, we investigate the influence of these aspects as independent factors on a user's understanding, perceived task load, and engagement. ***Contribution:*** *We show that making the user part of a data physicalization as an active element increases engagement and leverages perception beyond separate sensory channels.*

1.3 Dissertation Overview

The work presented in this dissertation comprises research work I did in the process of my doctoral studies at TU Wien. Each project resulted in a scientific publication, which I present in four separate chapters. Notably, P_2 received the "*Best Full Paper*" award at Eurovis 2025. These publications were co-authored with various colleagues, but always with me as the first author and main contributor. Chapters 3 – 6 are based on the following papers:

P_1 Pahr, D., Ehlers, H., Filipov, V. (2025). **HoloGraphs: An Interactive Physicalization for Dynamic Graphs**. In *Proceedings of the 20th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - IVAPP 2025, pages 859-866*. DOI: 10.5220/0013116000003912

P_2 Pahr, D., Di Bartolomeo, S., Ehlers, H., Filipov, V., Stoiber, C., Aigner, W., Wu, H.-Y., Raidou, R.G. (2025). **NODKANT: Exploring Constructive Network Physicalization.**, To appear in *Computer Graphics Forum 44 (3). Preprint at OSF*. DOI: https://doi.org/10.31219/osf.io/2wq97_v1.

P_3 Pahr, D., Piovarči, M., Wu, H.-Y., Raidou, R.G. (2024). **Squishicalization: Exploring Elastic Volume Physicalization.** In *IEEE Transactions on Visualization and Computer Graphics, vol. 31, no. 9, pp. 6437-6450*. DOI: 10.1109/TVCG.2024.3516481

P_4 Pahr, D., Ehlers, H., Wu, H.-Y., Waldner, M.,Raidou, R.G. (2024). **Investigating the Effect of Operation Mode and Manifestation on Physicalizations of Dynamic Processes**. In *Computer Graphics Forum 43 (3)*. DOI: https://doi.org/10.1111/cgf.15106

In addition to the publications mentioned above, I was also involved in the following publications as a co-author, responsible for mixed methods study design and data analysis:

- Ehlers, H., Pahr, D., Filipov, V., Wu, H.-Y., Raidou, R.G. (2024). **Me! Me! Me! Me! A study and comparison of ego network representations.** In *Computers & Graphics, Volume 125*. DOI: https://doi.org/10.1016/j.cag.2024.104123
- Ehlers, H., Pahr, D., Di Bartolomeo, S., Filipov, V., Wu, H.-Y., Raidou, R.G. (2025). **Wiggle! Wiggle! Wiggle! Visualizing uncertainty in node attributes in straight-line node-link diagrams using animated wiggliness.** In *Computers & Graphics, Volume 131*. DOI: https://doi.org/10.1016/j.cag.2025.104290.

Chapter 2 will summarize the background and related work relevant to this dissertation. In Chapter 3, I describe an interactive physicalization for dynamic graphs and investigate the benefits of natural vision on its readability. In Chapter 4, I present and study a toolkit for the creation of personalized network diagrams. I present an interactive physicalization pipeline for volumetric data, which I use to study elastic volume physicalization in Chapter 5. My study on the effects of manual operation of physical multisensory representations is shown in Chapter 6. Finally, in Chapter 7, I draw conclusions on the outcomes of the presented works and provide implications for future research in data physicalization.

CHAPTER 2

Related Work

Data physicalization is a multi-faceted, multi-disciplinary field within human-computer interaction (HCI) [DSMA+21, BFY+24]. In this section, I summarize the related work relevant to this dissertation, focusing on four core aspects. Firstly, I provide definitions and a general overview of data physicalization and describe works dealing with the concepts of interactivity and perceptional multimodality. Then, I present related works in three relevant fields of application for this dissertation: educational representations, as well as medical and network data physicalization. Moreover, I summarize related work in the context of fabrication and physicalization making. The chapter concludes with an overview of evaluation strategies for data physicalization.

2.1 Data Physicalization

While screen-based visualizations have replaced historical physical counterparts at the onset of the digital age, physical data representations remain most commonly associated with data-driven art. To break free from the "tyranny of the pixel," Andrew Vande Moere [Moe08, VMP10] first explores the potential of physicality in information visualization in 2008. In parallel, Zhao and Vande Moere [ZVM08] invoke physical embodiment as a means of creating metaphors to engage, especially laypeople, in an effort to democratize information, i.e., making the ever-growing collections of data more accessible to non-expert users. They define a data sculpture as a data-oriented physical form with artistic and functional purposes to augment an audience's understanding of the underlying data. Jansen et al. [JDI+15] coin the modern definition of data physicalization as encoding data into the properties of physical artifacts. This definition —purposefully— encompasses a vast design space, with the intent of describing non-virtual data representations of any kind. A highly related concept to physicalizations is multisensory representations, defined by Hogan and Hornecker [HH16] as a representation that seeks to reveal insight by encoding data into multiple modalities so that more than one sensory channel is required to fully interpret them. This definition is not intended to replace data physicalization but rather to characterize a specific way of encoding, and could also apply to virtual representations

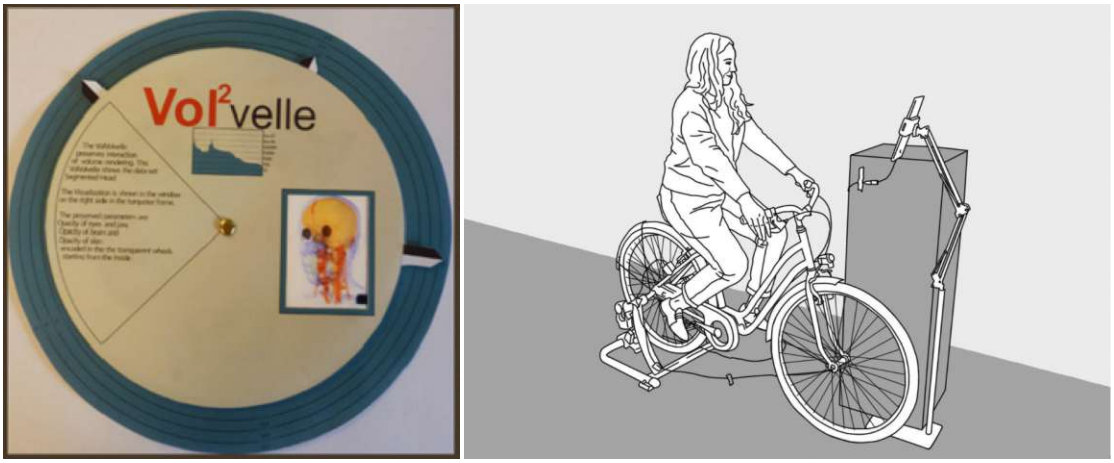

(a) (b)

Figure 2.1: Examples of data physicalizations with substantially different interaction and perception mechanisms. (a) Indirect interactivity: Rotating disks of transparent material allows users to customize a volume rendering, from Stoppel and Bruckner [SB17], © 2020, IEEE. (b) Direct interactivity, A treadmill to create the amount of energy used by a Google search, operated by a user, from Hurtienne et al. [HMC+20], © 2020, IEEE.

that are not artifact-based. Bae et al. [BZW+22] provides a description of the physicalization design space from different perspectives, providing many terms used in this dissertation. In conjunction with data physicalization, we often see the term *hands-on*. The applications of data physicalization presented in this dissertation deal with this aspect in different ways. Bae et al. [BZW+22] describe two distinct ways data physicalization can exhibit interactivity: directly and indirectly. An example of two physicalizations with indirect and direct interactivity is shown in Figure 2.1.

Indirect interaction constitutes manipulation of the representation similar to typical user interface operations [BM13] like filtering through data or selecting different visual parameters. Examples for this in the domain of data physicalization are fairly common. Stoppel and Bruckner [SB17] (Figure 2.1a) propose a system of rotating disks based on traditional rotating charts that lets a user customize a 2D volume rendering of medical data by altering the transfer function parameters. Similarly, Pahr et al. [PWR21] show a simple application of volume rendering by using volume data slices arranged in parallel at a fixed distance to create a 3D-like appearance. Individual slices can be removed and rearranged to investigate otherwise invisible structures. Schindler et al. [SWR20, SKRW22] propose the use of color filters to view separate layers of medical image data projected on papercraft. Using external electromechanical enhancements, Taher et al., [TJW+17] design a changing 3D bar chart that allows user interactions in real-time. Using the physical properties of conducting materials, Bae et al. [BFY+24] support users of their 3D network visualization pipeline with a method to detect when a node is selected by touch.

Direct interaction , on the other hand, goes beyond this by making the interaction part of the sensemaking process. Djavaherpour et al. [DMSJ23] show this at the hands of geospatial

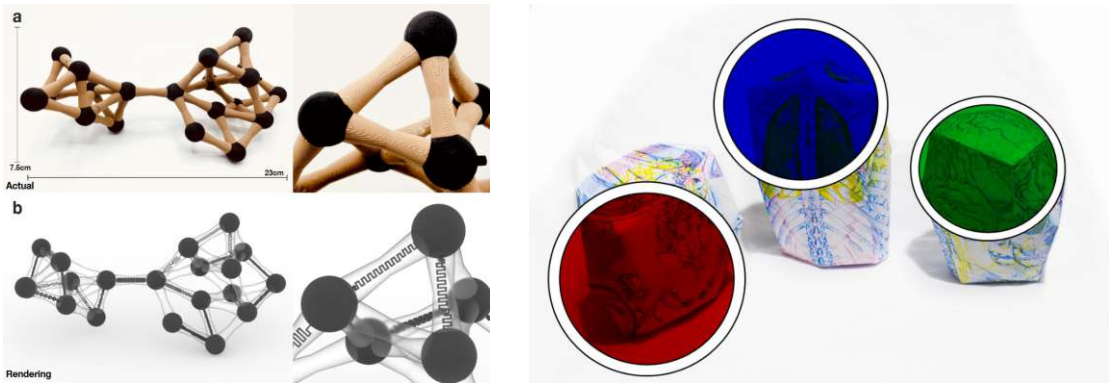

(a) (b)

Figure 2.2: Examples for different applications of data physicalizations presented in this dissertation. (a) A 3D interactive network physicalization that uses integrated electronic circuits for sensing, from Bae et al. [BFY+24], © 2024, IEEE. (b) An educational physicalization of medical data with multiple layers of anatomy printed on papercraft which can be viewed through color filters, from Schindler et al. [SWR20], © 2020, IEEE.

physicalization that encodes the roughness of arctic ice sheets into the surface properties of the model. Perovic et al. [PRC+23] study tactile interaction of users with data physicalizations using fluorescent markers that reveal touch traces using UV photography. Karyda et al. [KWK21] take this to a larger scale and propose three different person-sized physicalizations where users are presented with personal data for reflection, such as a foosball table, a treadmill, or a musical instrument. By changing the surface texture of an object, Ion et al. [IKS+18] encode continuous data into its tactile properties. A physical representation that completely involves the user in the information encoding is Hurtienne et al's. [HMC+20] "Move&Find" (Figure 2.1b). Here, a user has to create the energy required for a search query on a treadmill to communicate the "price" of using a search engine.

These examples show a clear relation between interactivity and multisensory representation. The studies presented in the course of this dissertation will highlight this interplay and further discuss their respective advantages and disadvantages.

2.2 Applications

Dragicevic et al. [DJVM20] define several areas of use for data physicalization. Specifically, the examples I present are concerned with either network data or volumetric data. In addition, I focus on educational applications for physical representations. Figure 2.2 shows an example of a physicalization from each domain.

Networks are an abstract representation of the relationship between entities. Visually, this is often represented by diagrams, such as node-link diagrams. Physical networks have been investigated with regard to the benefits of tactile interaction by Drogemuller et al. [DCWT19, DCW+21].

They test their hypotheses using 3D printed 2D node-link diagrams with solid edges. McGuffin et al. [MSF24] investigate 3D physical network diagrams as tactile props for augmented reality applications. Bar et al. [BFY+24] (Figure 2.2a) present a fabrication pipeline for interactive 3D network diagrams. These examples of network physicalization deal exclusively with static representations. To date, there is little research on the benefits of both indirect and direct interaction on data physicalizations.

Medical volume data is the product of 3D imaging workflows such as computed tomography (CT). Such data describe a three-dimensional structure, thus physicalization is often used to leverage natural perception. Ang et al. [ASS+19] study physical representations of blood flow models in an anatomical context by embedding different ways to represent flow directions into a model of the heart. Stoppel and Brucker [SB17] argue that their 2D volume physicalization serves the communication between medical professionals and laypeople. "Slice and Dice" by Raidou et al. [RGW20] is a method to create three-dimensional volume data sculptures using slices through volumetric data printed on orthogonally arranged sheets of transparent paper. Similarly, Pahr et al. [PWR21] use transparent material to create a holographic appearance from segmented volume data sets. Using foldable papercraft, Schindler et al. [SWR20, SKRW22] (Figure 2.2b) demonstrate an anatomical edutainment application that displays three-dimensional anatomy. Norooz et al. [NMJ+15] present a wearable anatomy puzzle in the form of a vest with detachable organs intended for the education of children. These examples communicate data exclusively over the visual channel. This demonstrates a lack of research on how data can be perceived multi-modally, for example, by adding tactility as a way to perceive sub-surface structures.

A core application for data physicalization is communication and *education*. Here, the purpose is not to analyze data, but to communicate important facts through the representation. In a recent literature review, Rau [Rau20] highlights gaps in the comparison of learning with virtual and physical representations. They indicated potential in researching the benefit of added tactile feedback to a physical representation and suggested isolating mechanisms in models to investigate this. Furthermore, they recommended accounting for cognitive load effects in the research of physical representations and the investigation of the impact of physical engagement on learning outcomes. Han [Han13] use the example of gearbox transmission to evaluate the influence of physicality on learning. Using a Lego model and a computer simulation, they investigate whether either form of manifestation would help understand the concept better. While neither model provided a significant advantage, they found that participants who had used a car with a manual gearbox performed better in the experiment. This could indicate a benefit of tactile feedback on the learning process. The examples above constitute physical representations of abstract concepts. However, in the context of dynamic physiological processes, such abstract physical representations are little explored.

2.3 Physicalization Making

Djavaherpour et al. [DSMA+21] provides an overview of fabrication and animation techniques for data physicalization. In this section, I present examples of physicalization toolkits and several fabrication pipelines. Later, I describe several techniques for the applications of 3D printing.

(a) (b)

Figure 2.3: Examples of different ways to create physical data representations. (a) A user creating a physical visualization with tangible tokens, from Huron et al. [HJC14], © 2014, IEEE. (b) A physicalization toolkit used to create accessible educational representations of volume data, from Pahr et al. [PWR21].

Figure 2.3 shows examples of how data physicalizations can be created.

2.3.1 Physicalization Toolkits

While there are numerous frameworks to create data visualizations for different types of data, the large design space for data physicalization does not allow the easy development of a general framework. As an exception, *MakerVis* [SSJ+14] supports the design of physical versions of well-known visualizations from data selection to fabrication. Conversely, in the context of volume data physicalization, several fabrication workflows and pipelines for the computer-aided generation of physical representations have been presented, targeting mostly laypeople edutainment. Schindler et al.'s [SWR20, SKRW22] workflow creates 3D (nested) papercrafts from digital meshes of anatomical or biological data. However, the underlying volumetric data have to be transformed into geometry meshes in a preprocessing step. Raidou et al.'s [RGW20] *Slice&Dice* demonstrates a pipeline for the computer-aided generation of 3D sculptures from volumetric data. They use regular printers to process transparent foils and octree-based spatial arrangements to create a 3D effect. Pahr et al. [PWR21] use similar materials and configurations for segmented medical volumetric data to allow the 2D/3D exploration of anatomical structures (Figure 2.3b). While accessible physicalization pipelines have been explored in educational contexts, there are few examples of fabrication methods to increase sense-making by supporting spatial perception.

2.3.2 Constructive Physicalization

In the field of *constructive visualization*, research has focused on democratizing the creation of visualizations, allowing non-experts to engage directly with datasets. Huron et al. [HCT+14] introduces a paradigm for users to create dynamic visualizations using wooden physical building blocks, shown in Figure 2.3a. They further demonstrate that constructive visualization enables novices to create visualizations [HCT+14] and spend more time on data-related tasks than using Excel [Kir10]. Recent work has developed several physical construction toolkits, such as Physikit [HGG+16], DataChest [WBH24], and SensorBricks [BVKVH24], which all use tangible elements, i.e., 3D printed tokens to create and explore data visualizations. Studies

have demonstrated the pedagogical potential of constructive visualization in both workshops and educational settings [HCBF16, KSB+23, WH16, Pun02, NP16]. These findings align with constructivist learning theory, where learners benefit from a "hands-on" approach through direct interaction with tangible objects rather than abstract methods [HJC14, VMP10]—a concept known as *discovery learning* [Pap93]. In LEGO® Serious Play [GV16], this concept is referred to as "thinking with your hands". Grounded in constructivist theory, it presents a method to facilitate problem-solving and communication, while creating an open system with infinite possibilities through creative play. The building is quick and straightforward and can be easily constructed, deconstructed, reviewed, and changed [Gau14]. While the examples above focus on people freely constructing visualizations, the impacts of construction on specific visual representations, such as graphs, warrant more research.

2.3.3 3D Printing Physicalization

Djavaherpour et al. [DSMA+21] identify additive fabrication as a suitable method for non-expert use to create detailed representations of limited size. Because of its accessibility, cost-effectiveness, and ease of use [GYM+17], several physicalization concepts have been fabricated using fused filament fabrication (FFF). For instance, Stusak et al. [STS+14] investigate different methods to create exercise feedback using 3D-printed artifacts. Similarly, Khot et al. [KHM14] propose 3D printed physicalization concepts for increasing the understanding of physical activity. Using *Fantibles* [KAL+16], the social media usage of sports fans has also been embodied in compact artifacts. In these cases, FFF printing is employed to create physicalizations for immediate feedback, supported by the relatively fast printing times of the method.

Additionally, FFF printers often offer simple modifications to create more complex prints. Ang et al. [ASS+19] use multi-nozzle FFF printing to create three-dimensional representations of blood flow in the human heart. Bae et al. [BFY+24] use multi-material printing to embed conductive material into their network physicalization to detect physical interaction on a microcontroller. Instead of using electronics, *AirTouch* [TRBA20] uses pneumatic sensing in elastic objects to detect user interaction. Also using pneumatics, Kim et al. [KET+21] present a pipeline to create inflatable elastic 3D-printed shape-changing objects. Torres et al. [TCKP15] introduce a fabrication toolkit, *HapticPrint*, to design haptically different interactive objects. Among other things, they use infill patterns to create different levels of elastic compliance. Fabrication devices typically offer only a limited selection of materials [SARW+15].

To enhance the design space, *metamaterials* are often used [BSP19]. These leverage the high spatial resolution of the fabrication hardware to generate tiny microstructures. Bickel et al. [BBO+10] propose to control the deformation behavior of 3D printed objects by using microstructures. In this way, it is possible to manifest bulk behavior similar to materials not available for printing. For instance, *Built-to-last* [LSZ+14] uses a honeycomb tesselation of 3D objects based on internal force distribution to maximize the strength-to-weight ratio in solid objects. Also, Martinez et al. [MHSL18] adapt Voronoi diagrams to the context of fabrication by replacing the distance function. They propose a special family of distance functions that guarantee fabrication constraints. Unfortunately, achieving control over the deposition at the micro-scale requires a tight coupling of the microstructure generation with the fabrication hardware. To this end, specific techniques have

been devised for inkjet printers [ZSCM17], powder sintering machines [SBR+15], photosensitive resin printers [MDL16], and fused filament fabrication [TTZ+20]. The examples above show the potential of consumer-grade FFF printing for data physicalization. However, metamaterial techniques have not been investigated in this context, specifically how they can be used to encode data attributes in the tactile properties of artifacts.

2.4 Evaluating Data Physicalizations

Evaluation is an inherent component of all aforementioned works. In the field of data physicalization and multisensory data displays, common comparisons involve virtual and physical representations [JDF13,SSB15,PMW+18,ASS+19]. Jansen et al. [JDF13] compare a physical 3D bar chart, virtual 2D, and 3D representations, finding that while the physical artifact outperforms the 3D virtual version, the 2D visualization is most efficient. Stusak et al. [SSB15] assess the memorability of physical vs. virtual bar charts, revealing better immediate memorization for the virtual chart but prolonged recall with the physical one. Ang et al. [ASS+19] compare a physicalized blood flow representation to a screen-based visualization, noting participants' faster interaction with the physical version. However, they find no clear evidence of superiority. The evaluation of "Move&Find" [HMC+20] focuses on understanding, engagement, and behavioral change, revealing that the multisensory representation enhances data understanding, creativity, and engagement. Drogemuller et al. [DCW+21] study the understanding of physical network representations. They compare physical representations both with and without tactile interaction, as well as virtual and haptic-only conditions. Their findings indicate heightened engagement, as well as self-perceived efficiency increase when people interact tactually with graph physicalizations.

Pollalis et al. [PMW+18] evaluate the usefulness of 3D printed artifact replicas compared with virtual ones on a screen and an AR device in an archaeology education setting. They use a mixed methods approach, measuring task times, enjoyment, perceived task workload, spatial presence, and learning outcomes both qualitatively and quantitatively. Their study highlights shortcomings of 3D printed artifacts due to inaccurate reproduction during the printing process. Taher et al. [TJW+17] and Sturdee et al. [SKA23] study interactions with an interactive physical bar chart in single [TJW+17] and co-located user [SKA23] scenarios respectively. Their findings show the great potential of interactive data physicalizations to encourage exploration. The co-located study also indicates positive social engagement when groups interact with a physical display together. Sereno et al. [SGBI22] compare different tangible interaction methods for the selection of data points in 3D space. Their findings indicate that AR devices have a lesser impact on subjective workload compared to screen-based methods. However, while users reported the interactions in AR to be more direct, this did not influence their performance.

Ferguson and Hegarty [FH95]compare learning effects achieved with physical representations of a mechanical system to diagrams abstracting the underlying concepts. They find that participants who were provided physical models are able to solve application tasks more accurately. Similarly, Boucheix and Schneider [BS09] compare static and animated learning aides. They show a positive influence of animations compared to static data representations for dynamic systems. Yet, the influence of user control on comprehension is not indicated. Finally, Gire et al. [GCC+10]

investigate the difference between physical and virtual learning aides. A virtual simulation and a physical model are employed in a comparative study to determine whether either type of manifestation provided advantages. They identify that certain aspects of learning are improved by physical interaction. Studies on data physicalization are typically targeted to evaluate the efficiency of a representation, especially compared to virtual counterparts. However, the emotional value of a representation is often neglected.

CHAPTER 3

HoloGraphs: An interactive Physicalization for Dynamic Networks

This chapter is based on the following publication:

> Pahr, D., Ehlers, H., Filipov, V. (2025). **HoloGraphs: An Interactive Physicalization for Dynamic Graphs**. In *Proceedings of the 20th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - IVAPP 2025, pages 859-866*. DOI: 10.5220/0013116000003912

Existing physicalizations of networks often focus on simple, static structures. However, dynamic networks introduce an additional layer of complexity, considering the temporal evolution of the entities and their relationships — an aspect that is crucial for understanding behavior over time [APS14]. Approaches typically tackle this problem by aggregating (or time-slicing) the network's temporal dimension [APP11, BDA+17], while continuous (event-based) representations capture changes occurring at finer temporal granularities [SAK20, AMA22, FCAA23]. We propose a fabrication pipeline that does not rely on 3D printing, resulting in an affordable and accessible physical data representation of dynamic networks, so-called "HoloGraphs"[1]. This modular representation affords interaction by allowing users to rearrange its pre-fabricated parts (Figure 1.2a). Here, the purely visual sensing always takes place after the physical action.

3.1 Introduction

Dynamic network visualization aims to support extracting insights and making sense of dynamically changing network structures [BBDW17]. Subsequently, dynamic graphs and their visual-

[1]The name "HoloGraphs" is an amalgam of the words "hologram" and "graph". In this context, it describes the technique to stack graph embeddings on transparent materials to create a 3-dimensional sculpture of a dynamic network.

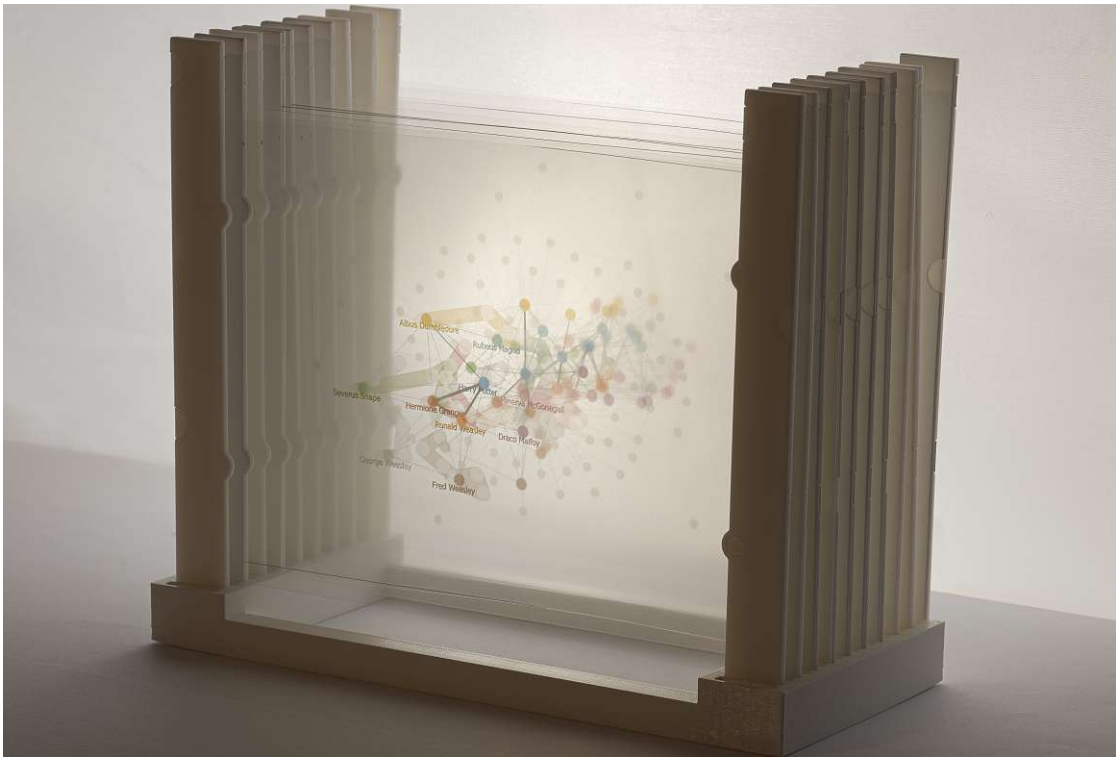

Figure 3.1: A *HoloGraph*. We display a dynamic graph by producing and embedding the individual timeslices, printing them on transparent media, and arranging them equally spaced. An overlay shows interesting nodes' trajectories over time and per individual timeslice.

ization have found common use across a variety of domains, from the social sciences [OKK13], through software engineering [RM14], to metabolic pathway analysis [RUK+10]. Several approaches to visualizing such graphs have been proposed, such as animated node-link diagrams [HMHU13], layered matrix representations [VBSW13], integrated 1.5*D* representations [SWW+15], or a superimposition of multiple 2*D* embeddings from different points in time [FCAA23]. With multiple embeddings for each point in time, we can produce *space-time cubes* ($2D+t$ or $3D$), that encode the temporal aspect in a third dimension. However, virtual 3*D* representations suffer from occlusions, distortions, and parallax effects [BPF14]. This can, in turn, obscure the graph's topology and change over time, negatively impacting user perception and understanding. In contrast, by arranging the slices side-by-side, i.e., juxtaposing them, we lose the ability to perceive nodes' movement over time between the individual timeslices [BBDW17, EMWR24, FABM24].

A key benefit of data physicalization is that we can examine 3*D* representations tangibly and intuitively by physically manipulating them. Research shows that certain drawbacks of 3*D* representations on screens can be overcome by such physicalizations [JDF13]. Before the use of computer graphics, the technique to stack panes of glass with electron density maps printed on was used to display density fields for molecules [Fox06]. Research regarding networks in

physical space has focused on both 2*D* embeddings with additional tactile encodings [DCW+21] as well as 3*D* embeddings with added interactivity [BFY+24]. Drogemuller et al. [DCW+21] present their findings about physicalized graphs being more engaging and fun to interact with, compared to their screen-based counterpart.

Given these outcomes regarding engagement and fun, it is unsurprising that interactive data physicalizations have value particularly related to education and visualization literacy [OSF04]. Manipulable physical representations have great value in educational settings, conceptually allowing for greater engagement, understanding, and learning [JDI+15,JD13,PEW+24]. We argue that there still are numerous opportunities to leverage the unique strengths of data physicalization to communicate complex scientific phenomena and concepts, such as dynamic graphs and space-time cubes.

We introduce a novel workflow to create interactive, physical representations of dynamic networks. We make this workflow accessible by moving the focus away from complex 3*D*-printing technologies and toward widely and cheaply available materials that can be assembled with various means. By printing network embeddings on overhead projector slides, we obtain transparent slices that can be assembled in parallel to create a 3*D* appearance: a *HoloGraph*. This allows us to keep the spatial perception of the *space-time cube* while enabling the examination of individual timeslices at any point. We divide each of the network's embeddings into *focus slices* that highlight and track nodes of interest and *context slices* that depict the contextual structure of the network at different points in time. We also provide separate global slices such as *label overlays* for nodes of interest and *trajectory overlays* to track nodes' movement over time. The codebase, printable versions of our networks, as well as 3*D*-printable meshes we used for our sculptures, can be found online[2]. We present a demonstration of our approach's utility in a case study, exploring the evolving relationships of characters across books of the "*Harry Potter*" series. In summary, the **contributions** of this work are two-fold:

- We present the development of a novel workflow with which to create physical and interactive representations of dynamic graphs.
- We demonstrate *HoloGraphs'* utility in a case study to highlight the value and importance of engagement when learning about concepts such as dynamic graphs.

3.2 HoloGraphs

The physicalization of node-link diagrams is conventionally done using 3*D* printing techniques [BFY+24]. However, such representations limit the number of nodes that can be printed and hence displayed owing to the time and cost associated with 3D printing; a limitation we here aim to side-step. Here, I [PWR21] propose methods to process and (interactively) display medical volumetric data using regular printers and transparent printable media. Moreover, outside of the context of data physicalization, Filipov et al. [FCAA23] describe a method to visualize dynamic networks by projecting the *space-time cube* embedding to a 2*D* representation. This

[2] https://osf.io/4u2e9/?view_only=751235378e564086beee9de8d37a6686

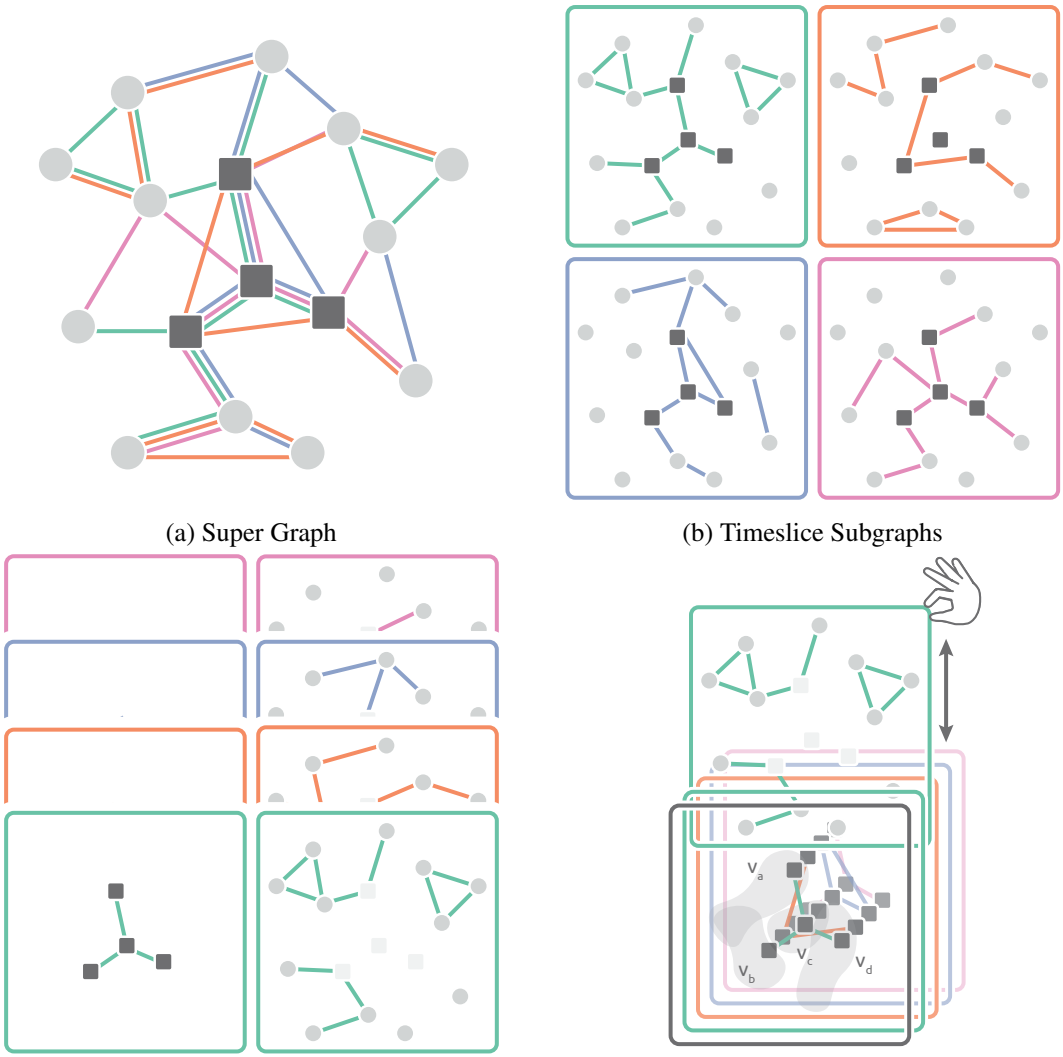

(a) Super Graph (b) Timeslice Subgraphs

(c) Focus+Context Subgraphs (d) HoloGraph

Figure 3.2: A dynamic graph (a), where connections between nodes and links differ between different points in time, is split up into timeslices t_1, t_2, t_3, t_4, representing the state of the network at different points in time (b). To emphasize nodes of interest (■), we divide the timeslices into focus (left) and context (right) subgraphs (c). Arranging the slices in parallel creates a *space-time cube* appearance (d). Individual timeslices can be removed for inspection and global overlays show the focus nodes' movements over time. For illustration purposes and simplicity, each timeslice subgraph shares the same layout. In practice, each timeslice subgraph is laid out semi-independently of the others, resulting in node movement between time points.

2*D* representation is an orthogonal projection of the network's topology over time in order to visualize node trajectories and behavior. Here, combining these two approaches, we propose to show the individual timeslices created by such a method in a 3*D* environment by printing each time-slice-subgraph embedding on transparent media and arranging them as parallel slices.

To add interactivity, we propose the use of various removeable **overlays** with which to display added **context**, draw **focus**, and provide **global characteristics** of the data. Figure 3.2 provides a simplified overview of the proposed process.

Definitions. Given a dynamic graph, formulated here as a time-sliced graph $G = (V_S, E_S, T)$, where V the total set of nodes and E the total set of undirected links across a set of timeslices T. For some timeslice $t \in T$ node-time-slice pairs (v_a, t) and (v_b, t) in V_S are connected by links $E_S \subseteq V_S \times V_S$, where $V_S \subseteq V \times T$. For a set of **focus nodes** $V_F \subseteq V$, we aim to highlight their position in the different timeslices over a disjoint set of **context nodes** $V_C \subseteq V$. First, the dynamic **super graph** G (Figure 3.2a) is separated into its constituent **time-sliced subgraphs** G_s (Figure 3.2b). These individual subgraphs are then each further broken up into their individual ***Focus+Context* subgraphs**, G_{sF} and G_{sC} respectively, such that $E_{sF} = V_F \times V_F$, as well as $E_{sC} = V_s \times V_C$ (Figure 3.2c). These *Focus+Context* subgraphs are then individually printed and mounted in a physical rack, forming (together with an additional labeled overlay that is printed along nodes' time-dependent trajectories) a tactile and interactive dynamic graph physicalization: the **HoloGraph** (Figure 3.2d).

Virtual Embedding. The **layout** of the network is computed using `D3.js`'s particle-based force-directed layout algorithm [BOH11]. Here, the key challenge of laying out a dynamic network is to strike a **balance between having enough change** in node placement between timeslices to effectively reflect the network's evolution over time, while also restricting said movement sufficiently in order to ensure **layout and network structure are still preserved**. Brandes et al. [BIM12] present several strategies for doing so, most notably *aggregation* and *anchoring*. In *HoloGraphs* we make use of an **anchoring** approach to computing the layout of the dynamic network: for each timeslice t, we utilize the layout of the previous timeslice $t-1$ as an initial layout before commencing the layout process. This ensures that the nodes' **movement over time will remain consistent** (i.e no flickering or popup effects) while preserving the viewer's mental map in the transitions from timeslice to timeslice [AP13].

Here, the movement of the nodes between timeslices depicts their trajectories over time as edges form or dissolve, pulling or pushing the nodes, respectively, reflecting the network's evolving structure. In contrast, an aggregation approach, i.e. computing singular positions for each node based on the layout of the super graph (Figure 3.2a), would result in no node movement over time. Figure 3.3a shows an example of the layouts of multiple timeslices overlaid on top of each other and highlighting nodes' movement over time (orange lines), making it difficult to make sense of changes over time. Instead, the focus nodes' trajectories, i.e. their movement between timeslices, can be better perceived in Figure 3.3c.

Physical Embedding. Figure 3.1 shows our version of a slide holder, used to stack the individual timeslices printed on transparent media. To make sure that our method supports a wide variety of fabrication techniques, we add **manufacturing parameters** to the physical embeddings. The physical representation depends on the **transparency** and **format** of the slides used. Additionally, the desired **colors** can be chosen depending on the availability of color printing.

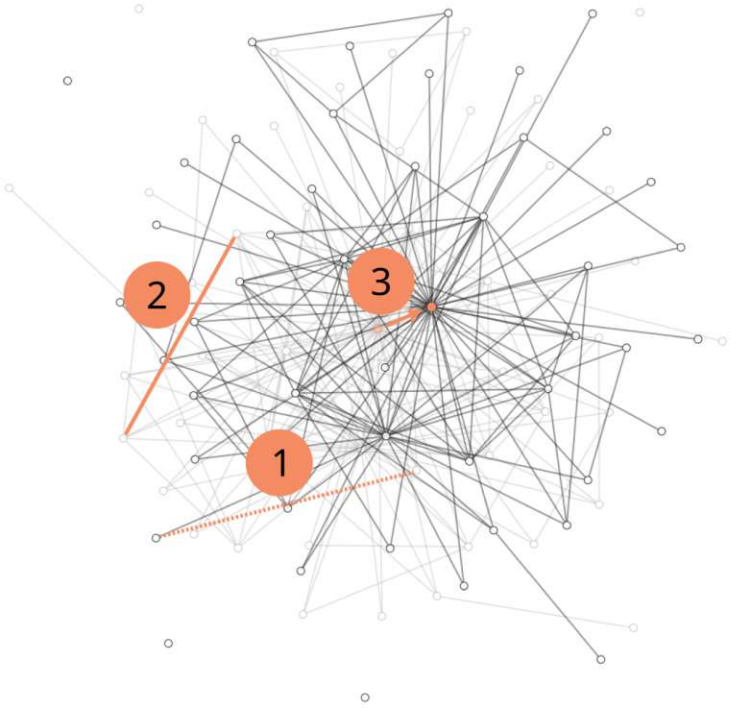

(a) Individual timeslices overlaid.

(b) Focus + Context

(c) Global overlays

Figure 3.3: Embeddings for individual timeslices, Focus + Context slices as well as Labels + Trajectories overlays. (a) The embedding of the subsequent slice is dependent on the previous. Disappearing (1) and appearing (2) links and the subsequent re-embedding causes movement of the nodes (3). (b) shows a superimposition of the focus subgraph, indicated by larger and colored nodes, and the context subgraph with smaller, faint grey nodes for a single timeslice. (c) shows the global overlays for focus node trajectories and labels.

First, we transform the node positions of the individual timeslices for the physical embedding to the format of the desired output. This corresponds to the **paper size**, with additional **margins** that ensure that the slide holders do not obscure any data. For all timeslices, we calculate the minimum and maximum x and y positions of the nodes. We then map the values to the space between the chosen margins. We then create **separate embeddings** for the chosen **focus** and **context** subgraphs. To emphasize nodes of interest, the focused subgraph is embedded using distinct node colors. In the context subgraph, we use smaller node diameters and a lighter hue of gray to avoid visual obstruction in the *HoloGraph*. The **edge weights** are encoded into edge thickness in both subgraphs. Figure 3.3b shows both of these embeddings superimposed. Printing the focus and context embeddings separately allows users to customize the representation interactively. Figure 3.4a shows the insertion or removal of a context slice from a *HoloGraph*. Finally, we also create **global overlays** to **trace the movement** of focus nodes over time, shown in Figure 3.3c. For this, we draw a polyline for each focus node through all its positions in the timeslices. At the node positions, we add a bright circle to emphasize where the node can be found. Figure 3.4b shows the process of adding this slice to the *HoloGraph*. We create a separate overlay for **node labels**. This way the labels can be removed freely, minimizing possible obstructions as shown in Figure 3.4c.

3.3 Case Study

To examine the potential of *HoloGraphs*, we present a case study highlighting social network interactions between characters in a famous children's book series. With this, we demonstrate the accessibility of our approach to people with little visualization literacy, and show that it reveals insights on dynamic data, despite its simplicity.

Data. We use a dynamic network representing the character interactions in the series of Harry Potter books by J.K. Rowling[3] to show the capabilities of our approach. The dataset is processed as a time-sliced network, where the nodes represent the characters in the books, and the weighted edges represent the number of interactions between those characters within each book. The individual timeslices represent the state of the character interaction graph within each book. We reduced the size of the original dataset to visualize the most central and influential actors as well as the most important relationships. By filtering out the links whose weights were not in the top 10 percentile we additionally reduce clutter and put an emphasis on the more important relationships between the characters of the data. This process resulted in a set of 111 nodes (characters) and 612 weighted edges (relationships) over 7 timeslices (books). From these 110 nodes, ten focus nodes, i.e. "egos", were selected (Figure 3.3c).

Implementation. Our pipeline's virtual embedding was done using `d3.js` [BOH11] and its physical embedding using `qt/python`. For the focus nodes, we use the top 10 nodes with the highest degree centrality, present in each of the seven timeslices. We used laser-printer-compatible overhead slides as the medium for the slices. The slide holder and base were 3*D* printed in a

[3] `https://github.com/nikhil-ravi/harry-potter-interactions`

Prusa MK2S. To use the entire printer surface, we decided to print the slices on half of an A4 page, resulting in 2 slices per page. The slices were attached to the holders using a hole punch on both sides of each slide to keep them in place. The slice holders fit into the left and right sides of the base, ensuring slices remained parallel when mounted. Figure 3.1 shows a fully assembled *HoloGraph*, Figure 3.4 shows how different overlays are added or removed.

(a) Context slices

(b) Trajectory overlay

(c) Labels

Figure 3.4: Composition of a *HoloGraph*. Focus slices show the nodes of interest at every timeslice. Context slices can be added to each focus slice individually (a). Node trajectories of focus nodes can be added as a global overlay (b), together with focus node labels (c).

Procedure. We invited a colleague with little experience in network visualization, but very familiar with the story of the Harry Potter books, to participate in our case study as a domain expert. After briefly explaining the concept of *HoloGraph* to them, we asked them to freely explore the physicalization and discuss notable events, character moments, and group interactions visible in the data. The session took about one hour. We present the notable findings and issues identified by our domain expert below.

Findings. Harry Potter, as the protagonist of all the books, takes a stable position in the center of the graph. *His movements are minimal; however, we can observe how other nodes of interest behave in relation to him.* Rubeus Hagrid, for example, is a central character in the first books, often interacting with Harry as his "window to the wizard world", however, in later books, he leaves the school and has other interactions. *While he remains tethered to Harry and in a stable orbit around him, his other interactions cause him to move away from the central set of characters*. Albus Dumbledore's interactions with Harry are often distant and occur at the end of each book. In books five and six, he becomes closely associated with Harry, and they form a personal relationship, culminating in private tuition sessions. *After his departure from the active cast in the sixth book, his node moves towards new characters related to his backstory and away from the central set of characters*. Fred and George Weasley are twins and are often shown interacting together. This is supported by their *very similar, sometimes parallel trajectories*. They *become separated close to the end of the last book when Fred's trajectory is a little shorter* than Georges's. Finally, one of the most faceted characters, Severus Snape also has an interesting trajectory. *He spends most of the time in a stable position, close to the center, as part of the "Hogwarts cluster"*. He also joins another faction dedicated to the protection of Harry in the later books, leading to continued proximity. *Finally, he changes factions again when his node is pulled away from the other teachers in the school and towards the antagonists*.

Issues. The node trajectories often point towards shifting associations of the characters with different groups. However, in the presented version of the Harry Potter *HoloGraph*, these groups are not highlighted. Our expert recommended *visualizing these groups by highlighting pertaining nodes in different colors in the context view*. Node trajectories could be more easily explained by having this added context. We also omitted different characters in favor of readability. Our selection of nodes of interest was purely based on node centrality. A suggestion here was to *emphasize a different selection of characters* that interact in different ways with the central characters. Notably, Voldemort, the primary antagonist of the story, was not included in the focus nodes for lack of centrality. Other nodes, such as the love interests of different characters, could have interesting trajectories in conjunction with the associated characters. Finally, while the trajectory overlay was overall helpful in finding nodes of interest in the different timeslices, it led to some confusion. Our expert thought that a certain move of the node occurred early on in the data, but noticed that the node was in a different place than expected at that time. The participant suggested *indicating a direction in the trajectory lines*, for example, by narrowing the lines from first to last timeslice.

3.4 Conclusions

Takeaways. Our work presents an **accessible** and **affordable** way to create **interactive** physicalizations of dynamic graphs. Our overarching goal is for this approach to be applied to improve visualization literacy and provide an engaging method for communicating science-related concepts in education settings. In our work, we investigated small to medium-sized graphs using the Harry Potter dataset. However, the number of nodes that can be displayed is limited only by printer resolution and the size of the used medium. Because we are printing our embeddings with 2*D* printing devices, the number of nodes can be much larger than in 3*D* printed graphs, where maximum node numbers of around 30 nodes are common [BFY+24, DCW+21]. Our flexible approach allows for the use of different sets of focus nodes as well, so that different aspects of the data can be examined with minimal reassembly, where, for 3*D* printed representations, printing times are often a limiting factor [DCW+21]. Most of a *HoloGraph* only has to be built once, and the various slices of a dynamic graph can be printed quickly and inexpensively. Compared to the approach of Bae et al. [BFY+24], *HoloGraphs* provides interactivity without the use of electronics or other augmentation. The case study confirms that our *HoloGraphs* can convey several facts about dynamic graph visualization at the hand of a simple example. A person with knowledge about the underlying data was able to quickly reason about important story events and character moments regarding several nodes of interest. Finally, our domain expert also pointed out interesting aspects our representation could display: highlighting different groups in context slides could explain node movement in conjunction with group association, while the trajectory overlay could be easily adjusted to help find nodes at different time points.

Limitations. The number of timeslices that a *HoloGraph* can display depends on several material factors. With the overhead projector slides we used, significant obstructions became apparent using about 20 slices (10 focus + 10 context slices). We used A5 pages (A4 halved) in landscape format for our embeddings. With a standard printer, sizes of up to an A4 page per timeslice are possible. Framing has to support larger slice formats well enough for the slices to remain parallel. Drogemuller et al. [DCW+21] show that users prefer tactile feedback in physical networks. Our approach, however, sacrifices the tactile component in favor of visual clarity.

Future Work. In this work, we investigated timesliced networks; however, many real-world networks are not discrete. In future work, we aim to investigate the use of *HoloGraphs* for continuous network physicalization. Our evaluation presents some insights into how our method can help with understanding the underlying data, visualization techniques, and what insights can be extracted from dynamic network structures and behaviors. However, the versatility and accessibility of our method could support other types of dynamic data, where discrete flattening in the z-axis can be applied, such as geospatial data.

CHAPTER 4

NODKANT: Exploring Constructive Network Physicalization

This chapter is based on the following publication:

> Pahr, D., Di Bartolomeo, S., Ehlers, H., Filipov, V., Stoiber, C., Aigner, W., Wu, H.-Y., Raidou, R.G. (2025). **NODKANT: Exploring Constructive Network Physicalization.**, In *Computer Graphics Forum 44 (3)*. DOI: https://doi.org/10.1111/cgf.70140.

While prior work has explored physical and digital constructive approaches for traditional data representations (e.g., line charts, bar charts), there is a clear gap in utilizing these approaches to comprehend the relationships and the patterns of underlying network data. Studies on how user-generated graph layouts differ from automatically generated ones in the virtual space exist [vHR08, DLF+09]. Yet there are unique considerations for constructive approaches to the physical space. Thus, we investigate how different construction approaches can engage people to overcome these challenges. We investigate the impact of constructive approaches on network comprehension, building on network physicalization research [DCW+21, MSF24, BFY+24] and insights from constructive visualization [BZP+20, HCT+14]. The tactile interaction occurring in the process of creating the network constitutes a first step towards the integration of user actions into the representation. The interactions during constructions allow the sensemaking process to begin before the representation is finished, but the tactile channel is not used for information encoding (Figure 1.2b).

4.1 Introduction

Network data comprise structured information that captures (complex) relationships, connections, and interactions between entities. Such data are encountered across various fields [RA15a], including social sciences, biology, and computer science [LPP+06, APS14]. Several representations—

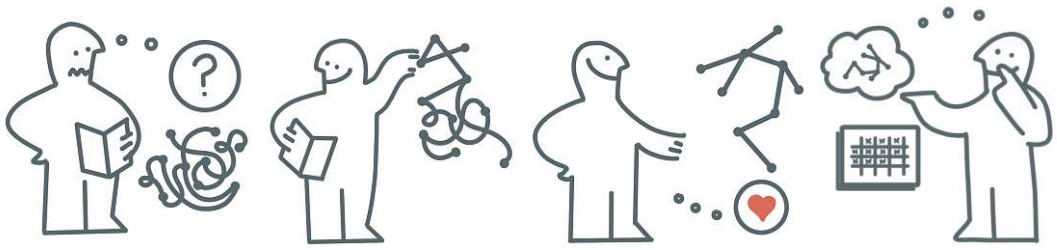

Figure 4.1: The constructive network physicalization pipeline: A user is interested in analyzing relational data. Using instructions, and the **NODKANT** toolkit, they construct a node-link diagram, generating insights about the data on the fly. The user then analyzes the physical diagram, leveraging its versatility and engaging with it interactively. Insights generated during construction can be recalled after some time.

ranging from traditional node-link diagrams to adjacency matrices, and to hybrid or alternative approaches—have been proposed for visualizing network data [FAM23]. However, many individuals struggle to interpret the physical meaning behind these representations, potentially due to limited familiarity with network visualization [BMBH16, ASSB+23, GTS10]. Unlike other visualization techniques, such as a bar chart or a scatterplot, network visualization requires the interpretation of complex underlying data relationships. This suggests that both local characteristics (i.e., connectivity, attributes, etc.) and global topological patterns (i.e., network density, hierarchical structures, clusters, etc.) need to be easily distinguishable. Recent research underscores that core concepts in network visualizations remain challenging for many, limiting their ability to derive meaningful insights [SCP+16, ASSB+23]. Overcoming these challenges requires new methods that promote a deeper engagement with the data.

Data physicalization refers to the process of transforming abstract data into tangible forms, allowing users to *interact* and *engage* with them physically [JDI+15]. Compared to traditional screen-based methods, this approach makes complex data more accessible by stimulating perception through direct interaction—fostering a deeper understanding and active engagement with the data. Conversely, constructive visualization (i.e., the construction of visualizations from physical tokens) encourages users to *build* data representations and stimulates deeper *reflection* on data compared to traditional tools [HCT+14]. This approach enhances learning and comprehension [Pap93] and also facilitates hands-on, personalized data representation [WBH24]. Similar principles are leveraged in products like IKEA® furniture or LEGO® [Gau14], where self-assembly creates a sense of accomplishment and enjoyment, experiencing the so-called "IKEA effect" [MNA12]. Combining data physicalization with constructivist principles offers an opportunity for network visualization to support users in building a mental map, while also enhancing their understanding of complex data relationships through physical engagement.

The combined potential of data physicalization [JDI+15] and constructive visualization [HCT+14, HJC14] has not been explored in the context of network visualization—particularly in terms of guiding users during the construction process. Specifically, whether users construct their networks freely or follow visual instructions, may significantly impact their efficiency, accuracy, and overall engagement with the data. To explore these factors, we developed a constructive network phys-

icalization toolkit, **NODKANT** (see Figure 4.1), which enables users to build their own network representations (specifically, node-link diagrams) with physical tokens. **NODKANT** is a playful reference to IKEA's product naming tradition, blending "nodes" and "links" to embody our constructivist network physicalization toolkit. We, then, investigate how the way data is presented influences users' construction and interaction processes, comprehension, and memorability throughout and after the construction process in a mixed-methods lab study. Our *contributions* comprise:

- **NODKANT**: A network construction kit consisting of a magnetic surface, 3D printable physical tokens for the nodes, and edges of adjustable length. This allows the construction of node-link diagrams with spools, representing nodes, and threads between them, representing the edges of the network (see Section 4.2).
- A mixed-methods lab study to assess the users' construction process, comprehension, interaction mechanisms, and memorability throughout and after construction using our kit (see Section 6.2).
- Our findings on the benefits and challenges of constructive network physicalization (see Section 4.4), supplemented by a discussion on implications for network physicalization (see Section 4.5).
- Open resources: All materials, data, code for generating instructions, and 3D printable meshes are available on osf.io.

4.2 Nodkant: A Network Physicalization Toolkit

We now set out to design a toolkit for network data physicalization. To begin with, we define a set of design requirements for such a network physicalization toolkit. Huron et al. [HCT+14] suggest that constructive visualization can profit from three creation paradigms: **simplicity, effectiveness, and dynamism**. We contextualize these paradigms in the scope of network visualization and our designs.

[Simplicity] ***Minimal number of parts and maximum amount of personalization.*** Using a simple case, Huron et al. [HJC14] demonstrate the versatility of square, colored, wooden tokens to create a multitude of different data representations. We follow this concept by minimizing the amount of unique parts in the representation, while simultaneously providing users with as much freedom in creating their personalized representations as possible. We argue that simple elements, when thoughtfully designed, can serve as versatile building blocks for complex systems. At the same time, they streamline the design, reduce cognitive load, and facilitate use.

[Effectiveness] ***Familiar and accurate representation of the concept.*** Networks have long been used by humans to represent entity relationships [MFD20], with node-link diagrams being arguably the most common representation [Tam13]. Applications like social network visualization [DLM24], transport networks [WNT+20], and biological networks [EBK+25] are intuitively understood methods to abstract complex phenomena.

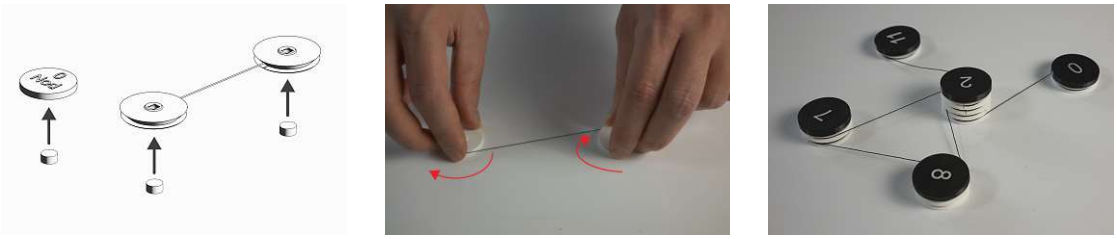

(a) (b) (c)

Figure 4.2: Steps to create a **NODKANT** diagram. (a) Nodes are 3D printed and magnets are placed underneath, while edges are fitted with a length of yarn. (b) Edge lengths are adjusted by turning the spools until the edge has the desired length. (c) Edge spools and nodes are stacked vertically on a magnetic surface, where the magnets ensure stability until, step by step, a physical network is formed.

[Dynamism] ***Make use of the physical nature of the representation.*** Data physicalizations possess unique abilities to engage audiences purely through their tangible nature [ZVM08, WSK+19]. Digital representations commonly provide benefits such as interactivity; however, tactile interactions have also been explored and show benefits for physical network representations [DCW+21]. Recent developments, such as Bae et al.'s [BFY+24] computational pipeline to incorporate sensors into data physicalizations, represent a step towards leveraging the potential of interactivity in data physicalizations as well [PEW+24].

Designing NODKANT. Graphs that use edges represented as lines can benefit from improved readability [ANMMG25]. Straight-line node-link diagrams represent only one way to visualize network data [Tam13] and comprise two fundamental parts—*nodes* and *links*. In these diagrams, nodes represent entities or objects, while edges represent their relationships or connections, visually forming a network. These **simple** components can be used to easily create diverse network representations.

With **NODKANT**, we intend to provide an **effective** strategy to represent networks—akin to the way they are represented in literature and practical applications. Bae et al. [BFY+24] use spheres to represent nodes in their 3*D* network physicalizations. Their pipeline focuses on the sensing capabilities of their model and does label nodes to make them identifiable. Conversely, McGuffin et al. [MSF24] emboss the node labels into cuboids to represent nodes in their 3*D* network models. Drogemuller et al. [DCW+21] use spheres to represent nodes. Due to the limited size of their model, they also omit placing node labels. In **NODKANT**, a node is embodied by a spool, i.e., a flat cylinder as shown in Figure 4.2a. This shape can be easily produced using digital fabrication tools like a 3*D* printer or laser cutter. The top of the cylinder accommodates space for a node label, which can be simply written, attached as a sticker, printed, or laser-engraved.

Finally, we support the **dynamic** creation of our network physicalizations. When using fabrication to create network physicalization, creating nodes and edges in a single step limits the materials that can be used. The existing network physicalization methods [BFY+24, MSF24, DCW+21] create solid edges between the nodes. While this consolidates the network's structure, it limits the size of the representation due to the print surface and constrains the corresponding interactivity.

To support the step-by-step creation of physical networks, we propose to use yarn to represent edges (Figure 4.2a). Using digital fabrication tools, we create a spool attached to each end of a length of yarn. This enables users to easily manipulate edge lengths if needed (Figure 4.2b). Placing a magnet underneath the spool creates a rotational axis for the spool at the contact point, which allows the edge lengths to be adjusted after the spool is placed. We design the spools to stack on top of each other, allowing a single node to have multiple connections (Figure 4.2c). The magnets provide enough stability to stay connected while the resulting stacks can be easily moved around on the surface. Note also that, based on this setting, a user can loosen the yarn to form various edge styles other than straight lines depending on their preferences.

4.3 Studying Constructive Network Physicalization

We assess our proposed **NODKANT** toolkit in an exploratory study evaluating the influence of personal construction on a user's comprehension of network data. The theory of *embodied cognition* in cognitive psychology proposes that human perception is strongly connected to our interactions with the environment and how this influences learning processes [NEFM99]. Suwa and Tversky [ST02] show that creating representations of concepts, such as sketches, facilitates idea generation. Huron et al. [HJC14] investigate this further by employing generic physical tokens with which users freely create personalized data representations. Inspired by these approaches, our approach constitutes a physical method to create external representations for an abstract concept. While we restrict the visual representation to node-link diagrams, users retain the flexibility to determine the *layout*, which is an important part of the presentation mapping (geometry) [MGWP23].

4.3.1 Research Questions

The potential benefits of constructing a personalized physicalization, combined with the unique challenges inherent to network visualization present an interesting opportunity. We design our exploratory study based on four distinct research questions:

[RQ1: Construction] *What is the influence of different ways to present the data on user experience, while* ***constructing*** *the networks?* Constructive visualization enables a user to create a personalized representation of a given dataset. Limiting the representation to a node-link diagram still leaves plenty of room for customization. We investigate the impact of different methods to support a user during construction [WHJ23], on the user's experience.

[RQ2: Comprehension] *How does constructing a network physicalization impact a user's* ***comprehension****?* When users construct a personalized node-link diagram, they embody part of the presentation mapping in the visualization pipeline [JD13]. This process enables user interaction—even before the physical representation is fully realized [HJC14]. Conversely, user-created layouts may hinder readability compared to algorithmically optimized designs. To investigate this, we measure users' insight generation and performance in network analysis tasks, depending on *whether* and *how* they constructed their physicalization.

[RQ3: Interaction] *Which **interaction** patterns emerge from the use of **NODKANT**?* With **NODKANT**, we propose a network physicalization toolkit that supports the creation of node-link diagrams. Physicalizations inherently afford physical interactions through their physical embodiment [JDF13]. In our case, we see opportunities for interaction both *during* and *after* the construction process. To investigate this, we monitor users' interactions and analyze behavioral differences based on their construction method.

[RQ4: Memorability] *Does constructing a network physicalization impact **memorability**?* Physicalizations are more memorable than screen-based visualizations [SSB15, HMC+20]. To investigate *why* data physicalization supports users in remembering data, we measure delayed insight generation and task performance in a follow-up study.

4.3.2 Metrics

Quantitative Metrics. During the construction session, we measure the **time** taken by the participants during construction. To assess their subjective experience, we use the **NASA-TLX** questionnaire [HS88] as a benchmark. Additionally, we measure users' **time** and **accuracy** in completing benchmark tasks. For efficiency, we omit task-specific NASA-TLX evaluations and instead collect a single **subjective difficulty rating** from participants.

In the follow-up study, we ask participants **closed questions** about global properties like the represented entities and study context. We ask participants to recall known numeric properties and derive properties they did not calculate before. Finally, we include a Likert scale rating section into the follow-up to gather participants' subjective assessments of **NODKANT** by using four questions that relate to the **value** of visualization according to Stasko [Sta14] ("*The presented visualization was... i) trustworthy, ii) understandable, iii) available, and iv) quickly accessible*") and five questions that refer to the **emotional value** of visualization ("*The presented engaged in a... i) creative, ii) affective, iii) intellectual, iv) social, v) physical way*"), as defined by Wang et al. [WSK+19].

Qualitative Metrics. We **record** the construction sessions with an overhead camera setup and think-aloud protocol to analyze construction strategies, as well as interactions and thoughts during benchmark tasks. After construction, as well as the individual benchmark tasks, we ask participants to summarize their thoughts during construction and transcribe this as **open feedback**. Before the benchmark tasks, we record the open exploration of the network by the participants for an **insight**-based evaluation. In the follow-up study, we ask **open questions** about global and local structures and analyze the answers for insights.

4.3.3 Conditions

Our toolkit enables users to create personalized embeddings of network data, as we assume that the way **the dataset is presented** to the user may have a significant impact on their comprehension. Existing studies on user-created layouts in digital environments [vHR08, DLF+09] confronted users with unstructured layouts to rearrange or draw a layout using an edge list [PPP12, Pur14]. Our study focuses on the physical creation of network diagrams, and the associated benefits and

pitfalls. Wei et al. [WHJ23] present different strategies to guide users in creating personalized visualizations: *next-step*, *ghost*, and *gallery*. For our use case, we present a fixed visual mapping, the node-link diagram; thus, we will not investigate the gallery-based option, where users are shown representation alternatives to choose from. We base our conditions on these techniques.

Free Construction (FC). In *next-step*, a user is shown possible positions for the next step to take. In our case, this can be easily translated into a step-by-step placement of edges (Figure 4.3, FC). Edge lists represent network data by listing each connection between two nodes. When working with a randomized edge list, users can reconstruct the network by placing edges one at a time. However, this process often requires extensive rearrangement due to the randomness. We, therefore, opt to present the data to the user as a **sorted edge list**. Different sorting criteria are possible for this [Mey79], and to choose the most suitable one for our case, we compared sorting edges by *degree* of associated nodes, and following a *spanning tree* in a pilot study (see Section 4.3.4).

Layout Construction (LC). *Ghost* shows an outline of the final representation to the user, to indicate where tokens should be placed. The spatial embedding of a network is an area of ongoing research. To assist users in this process, we propose to provide users with a **pre-computed layout** as a means of presentation (Figure 4.3, LC). A force-directed layout algorithm ensures that the resulting layout conforms to established aesthetic criteria [Tam13]. To prevent users from getting lost while constructing their network, we provide step-by-step instructions on how to recreate the given layout, similar to instructions found in LEGO® or IKEA® manuals. To maintain consistency across conditions, the steps are presented in the same order as the edge-list in the (FC) condition.

No Construction (NC). A **NODKANT** diagram is a data physicalization—regardless of how (and by whom) it is constructed (Figure 4.3, NC). In the NC condition, we provide a group of participants with a **pre-constructed representation** of the layout, serving as a baseline for our comparisons.

Figure 4.3: Study conditions. FC: Users freely construct their network from a sorted edge list. LC: Users construct a pre-computed layout using a step-by-step manual. NC: Users solve network analysis tasks using a pre-constructed layout.

4.3.4 Pilot Study

As a first step, we conducted a small pilot study among six co-authors to determine which network size is suitable, i.e., the upper limit for a lab study. Additionally, as mentioned above, we compared two different sortings of the edge list for free construction scenarios to determine the most appropriate.

Procedure. We selected three networks (`grafo2693.13`, `grafo634.24`, and `grafo1034.29`) from Rome-Lib [DGL+97] based on node count, edge count, and density. Each participant assembled all three networks following a counter-balanced schema also alternating between **FC** and **LC**. For **FC** sessions, we tested two edge list sorting algorithms, one sorted by node degree and the other sorted along a spanning tree following node degree for non-spanning edges. We measured the time for each run and discussed emerging thoughts and implications.

Results. Completion times mainly depended on the edge number in the network, with the smallest network taking around 6–10 minutes, and the largest one taking 20–27 minutes to construct in **FC**. We noticed a broader variation of construction times in the **FC** condition compared to **LC**, owing to different strategies employed by the participants. In summary, we made the following decisions for the lab study based on our findings: Study times around 1 hour, we decided to use a dense **network with around 30 edges**. An edge list formed around a spanning tree leads to participants inefficiently using construction space sometimes, prompting us to favor the **sorting by node degree**. Dwyer et al. [DLF+09] suggest it may be intimidating to present users with a layouting task without usecase; therefore, we decided to **use data with relatable or recognizable context**.

4.3.5 Data

In our pilot study, we discussed the importance of a salient use case for our tasks. We decided to use the **animal contact network** from the network data repository [RA15b], created for studying rabies propagation [RHGC15]. In the network, nodes represent raccoons and edges represent recorded contact interactions between them. In our pilot study, we identified that a network with approximate 30 edges should be feasible to construct within 30 minutes. Thus, we selected the `mammalia-raccoon-proximity-50` network with 16 nodes and 33 edges and assigned randomized names as labels to make it more relatable. The network density (0.1375) provides a sufficient challenge for analysis tasks [YAD+18], while the low number of nodes limits task complexity. Also, the origin of the data provides context for users to interpret the tasks.

4.3.6 Tasks

We select a set of tasks for our user performance evaluation from the network task taxonomy by Lee et al. [LPP+06]. Before exposing participants to benchmark tasks to measure completion times and error rates depending on the experiment condition, we ask them to freely examine the network. This allows us to transcribe the resulting statements and measure insight generation as proposed by North [Nor06]. We follow this with numeric observations about the **general**

overview—namely counting or estimating the number of nodes and edges in the network. The remaining tasks are assigned according to a balanced scheme and we present them here in the order of occurrence in the taxonomy [LPP+06], orienting on *friend-of-a friend* scenarios. We select several topology-based tasks:

[Adjacency] *Find the most and least connected nodes in the dataset.* We use the context of popularity in a social structure and the corresponding risk of infection.
[Direct Accessibility] *Find out if three pairs of nodes are connected.* We employ the contexts of friendships and direct disease transmission.
[Indirect Accessibility] *Find all nodes accessible at a hop distance of two* for three different nodes. We employ the contexts of a *friend-of-a-friend* and the spread of infectious diseases.
[Common Connection] *Find the common neighbors* of three pairs of nodes. We employ the contexts of having common friends and infection by common contact.
[Connectivity] *Find bridge nodes in the network*, defined as nodes that upon removal cause a split of the network into components with at least two connected nodes. We use the contexts of connecting friend groups and isolating risk patients.
[Paths] *Find the shortest paths* between three pairs of nodes. We employ the contexts of connection in a social network and critical paths in contact tracing.

4.3.7 Recruiting

We recruited 27 university students, aged between 20 and 28 years for our study. Five of the students identified as female, the remaining 22 as male. Four of them had already completed a bachelor's program in computer science, and 23 had a high school diploma. All participants reported basic knowledge about networks, having completed a course on algorithms, where the concept is first presented. None of the participants had further education or professional experience with networks in a visualization context.

4.3.8 Procedure

We employed a between-subject design in our study. Participants were assigned to one of three groups, corresponding to one of the three conditions (**FC**, **LC**, **NC**) using a balanced scheme. All study participants received onboarding about node-link diagrams. We first explained to them the meaning behind the network, i.e., nodes representing individual animals, and the edges between them represent the interactions. We listed them two possible use cases for such networks: investigating i) social structures and ii) the corresponding ways of disease transmission. Every participant was instructed about the parts of **NODKANT** and was explicitly encouraged to interact with it freely during the tasks.

Construction. Participants from groups **FC** and **LC** received training on how to construct a network using **NODKANT**. They were given the option of free interaction with its components until they felt confident with them. After completing the scale ranking procedure of the NASA-TLX [HS88], participants were handed a printed set of instructions to construct the network: **FC** participants received a sorted edge list, while **LC** participants received an instruction booklet.

After construction, we asked participants about their thoughts during the process. Conversely, participants in the NC group received a pre-constructed representation to observe instead of completing this step themselves.

Benchmark Tasks. With their physical networks in hand, we asked participants to reason about interesting global structures and interesting nodes they could detect. First, we asked them to count or estimate the number of nodes and edges. Then, participants completed a set of network exploration tasks. We provided a context for each task for the two use cases (social structures and disease transmission), explaining how the task may be relevant in a social network and disease-monitoring scenario. The order of tasks was counter-balanced across the participants using a Latin square scheme. We asked participants to think aloud during the tasks and leave feedback after each task. After all tasks were completed, we requested them to comment on at least one positive and one negative experience during their experiment and asked for additional informal feedback. Finally, we thanked the participants and concluded the on-site experiment.

Follow Up. Each participant received an online questionnaire exactly 10 days after completing the on-site study. We asked the participants to complete the questionnaire within four days of receiving it. We sent two reminders, two and four days after sending out the questionnaire. All but three participants completed the questionnaire within four days, two responded after five days, and one participant submitted the response after nine days. Participants were first asked to recollect the context of the experiment and what nodes and edges were represented. We asked them to **freely recall** list interesting global and local structures in the network, without cues. Additionally, we provided them with **cued** questions on the number of nodes and edges, and edge density—a detail which they were not asked to calculate during the on-site part, to avoid receiving only memorized answers. Lastly, participants completed two Likert scale ratings about **NODKANT**. At the end of the questionnaire, participants were again thanked and received a debrief about the purpose and procedure of the study.

4.3.9 Analysis

Quantitative Analysis. Due to our between-subject study design, we obtain a sample of 9 participants per group, preventing meaningful statistical analysis. Thus, we omit statistical testing on the results and discuss the results purely visually. We compute the NASA-TLX [HS88] scores as weighted averages across the scales *Mental Demand, Physical Demand, Temporal Demand, Performance, Effort,* and *Frustration*, with weights as obtained in the scale ranking procedure. Task **accuracy** is computed as the ratio of correct observations to total observations in tasks with a set of nodes as answers, e.g. ground truth is (A,B), reported set is (A), accuracy is 50%. When the task is to determine a number, we report the accuracy as $1-|e_{rel}|$, where e_{rel} denotes the error relative to the ground truth, e.g. ground truth is 10, reported number is 9, $e_{rel}=10\%$, and accuracy is 90%. For tasks that have three sub-tasks, we report the average accuracy across sub-tasks. We classify the **answers to the closed questions** in the follow-up as correct, semi-correct, and incorrect. For example, an answer to "*Which entity did the nodes represent?*" could be "*raccoon*" (correct, 100%), "*animals*" (semi-correct, 50%), or "*people*"

(incorrect, 0%). We report the subjective **task difficulty** on a 1–10 scale (1-easy; 10-difficulty). Task and construction **times** are reported in m:ss. The time runs from when participants receive the assignment until they confirm to be satisfied with the result. In the case of subtasks, we report the total time for all tasks. In the subjective Likert scale ratings, we asked participants to rate their experience according to the scale *Disagree* → *Somewhat Disagree* → *Neutral* → *Somewhat Agree* → *Agree*.

Qualitative Analysis. We further **analyzed video** recordings and categorized the degree of interactions of our participants. We distinguish three categories: (1) the participant is *passive* or simply points at the representation, (2) the participant *touches* the network during the session without disturbing the layout, and (3) the participant *moves* the parts around.

We transcribed the participants' open feedback and further processed it, resulting in 521 individual utterances. We conducted two rounds of coding, one for categorization and one for rating insight-related utterances resulting from the first coding. Both were conducted by three independent coders (one was present in both procedures). For the **categorization**, we follow a combination of inductive and deductive coding [Chi97]. In the inductive session, each coder independently assigned a single concept to each utterance. Each coder produced their own set of codes, ranging between 21 and 105 different concepts. The coders then met in a joint session to produce a unified codebook resulting in 35 codes pertaining to 8 categories. After this, a deductive coding session took place, where the coders independently assigned these codes to the utterances. In another joint session, the coders discussed their assignments in cases where no consensus was reached. In a final individual session, each coder reviewed dissenting assignments. After that, the remaining 2-to-1 conflicts were resolved following a majority vote.

During the first coding procedure, we identified 149 comments from 5 different codes as **insights**. We categorize insights in three levels: (1) *reading data*, (2) *reading between data*, and (3) *reading beyond data* [SGA22]. The coders first assigned their rating. In a joint session, they discussed dissenting selections. In a second individual session, the coders revised their choices until no more conflicts remained that could not be broken by a majority vote.

4.4 Results

We present our results as they relate to each of our research questions (see Section 4.3.1), and summarize the implications of all relevant metrics (Table 6.1) to provide succinct answers.

4.4.1 RQ1: Construction

Quantitative Results. For **FC**, the median **construction time** was just under 27 minutes. The median for **LC** was about seven minutes faster. In the weighted average **TLX-score**, we observe a **higher median for FC** than for **LC** (**FC**: 33.67, **LC**: 22.67). We show the results for responses for the NASA-TLX across all sub-scales, as well as the weighted average in Figure 4.5. Notably, we observe a higher median in *physical demand* for **FC** (30) than for **LC** (20), as well as three very high ratings (≥ 80). We also observe higher medians in **FC** compared to **LC** for *performance*

Table 4.1: Relation of metrics to research questions. Conditions are FC (free construction), LC (layout constrcution), and NC (no construction). ↑ placed after a condition indicates that the metric had positive implications for it, ↓ means negative. ↔ denotes ambiguous implications between conditions.

		RQ1	RQ2	RQ3	RQ4
Quantitative	Construction Time	FC ↓			
	Construction TLX	FC ↓			
	Task Time		FC ↔ LC ↔ NC		
	Task Accuracy		FC ↔ LC ↔ NC		
	Task Difficulty		FC ↔ LC ↔ NC		
	Value (Traditional)		FC ↓		
	Value (Emotional)		FC ↔ LC ↔ NC		
	Follow-up Accuracy		NC ↓		NC ↓
Qualitative	Construction Strategy	FC ↑			
	Physical Properties	FC ↓ LC ↓		FC ↑ LC ↑	
	Interaction			FC ↑	
	Usability		FC ↓ LC ↓		
	Task Load		FC ↓ LC ↓		
	Engagement		FC ↑ LC ↑		
	Feedback	FC ↔ LC		FC ↔ LC	
	Insights		FC ↑		NC ↓
	Video			FC ↑	

(FC: 40, LC: 15), effort (FC: 30, LC: 20) and *frustration* (FC: 25, LC: 10). Interestingly, *temporal demand* shows a very similar distribution for FC and LC, while we observe equal medians in *mental demand*, but a higher variation of ratings in LC.

Qualitative Results. We found a total of 28 (FC:20, LC:8) utterances referring to strategies during **construction**. Five FC participants *followed instructions* ("*worked from node to node in order they came up*") and used our edge list sorting, while four reported individual *sorting strategies* ("*looked at which names come up more often for the best start*"). Six FC participants started *forming clusters* during construction ("*noticed clusters forming so I kept them in separate areas*"). Five FC and three LC participants remarked on rearranging nodes during construction ("*Moved nodes around a lot during construction*", and "*should have checked [the] final outcome before, [and] had to rearrange a bit*", respectively). Only LC participants expressed strategies concerning *label placement* during construction. 10 comments on **task load** were related to strain during construction (FC:4, LC:6). In FC, two comments referred to *frustration*, one to increased *temporal effort*, and one to low *performance*. For LC, one comment referred to lowered while three to increased *mental demand*, and three comments referred to increased *temporal* and *physical effort*, each. Two comments represented **feedback** on possible *physical improvements* of spool handling. **Visual inspection** of construction results shows that FC led to more "messy" representations than LC.

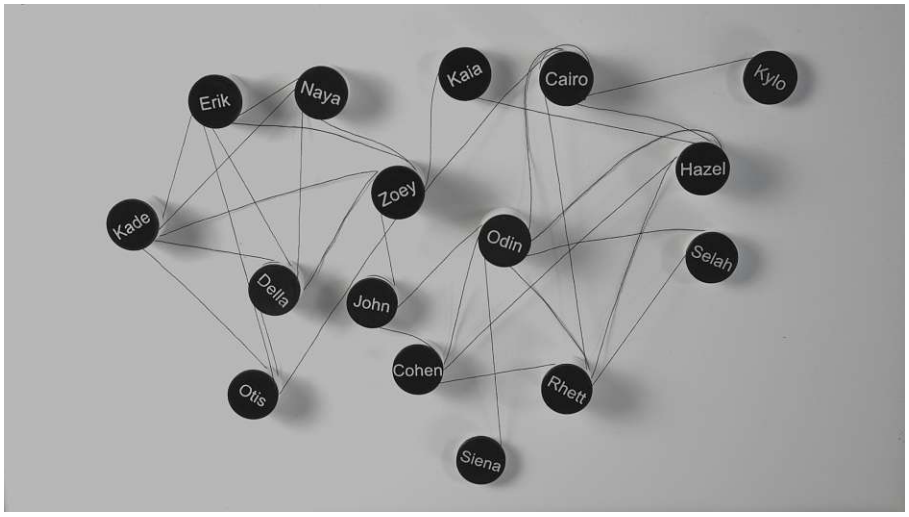

(a) FC layout

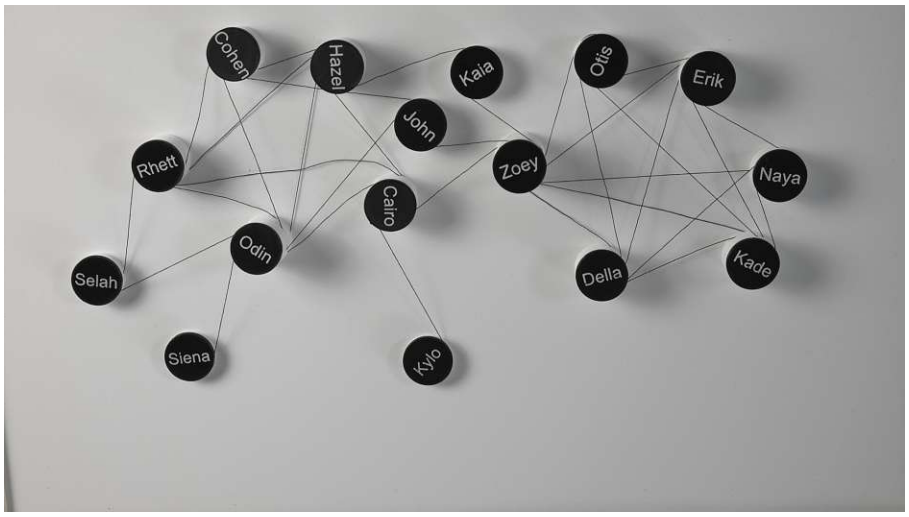

(b) LC layout

Figure 4.4: Example layouts. FC graphs use a lot of space and edges are often loose. LC graphs resemble the given layout, sometimes showing loose edges. NC participants were given a faithful recreation of the LC layout.

Summary RQ1: Construction. Our findings indicate that constructing diagrams with **NODKANT** takes longer on average and has a slightly higher impact on a user's task load when using FC compared to LC. Interestingly, the **TLX scores do not show an impact of the longer construction times in FC**. FC results are less refined compared to the outcome of LC (see Figure 4.4-a,b). Many negative comments overall refer to frustration caused by difficulties with string tension. Yet, we find indications in the qualitative data that during construction FC **participants were deeply engaged** and worried more about global issues like emerging clusters, while LC

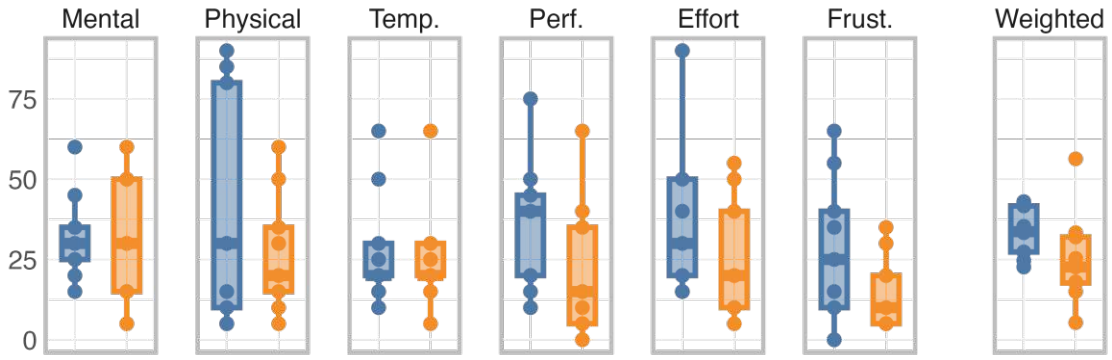

Figure 4.5: NASA-TLX ratings for network construction, blue plots relate to FC, orange to LC. FC participants report higher physical demand, overall effort, and frustration, as well as less performance satisfaction than LC participants, yielding a slightly higher weighted average score for FC.

participants' major concern lay in the aesthetic details, such as keeping the edges straight.

4.4.2 RQ2: Comprehension

Quantitative Results. Overall, task **completion times** between all groups were similar, with FC (*avg* = 01:52) being on average 30 seconds slower than LC and NC (both *avg* = 01:22). Because three participants took particularly long, the average **time for the *connectivity* task in LC was about double compared to LC and NC** (FC: 02:55, LC: 01:24, NC: 01:15). We observe notably higher median completion times for FC in the tasks *common connections* (FC: 01:46, LC: 01:15, NC: 01:27), *indirect accessibility* (FC: 04:08, LC: 03:20, NC: 03:12) and *overview* (FC: 01:10, LC: 00:40, NC: 00:41). We also observe high **accuracy** for all conditions on average across all tasks. In the *overview* task, we observe a lower median accuracy for FC (FC: 78.79, LC: 90.91, NC: 90.91). Across all benchmark tasks, LC participants rated the tasks **difficulty** highest on average, followed by FC, and NC (FC: 3.82, LC: 4.16, NC: 3.13). Median ratings for FC were notably higher compared to the others for the *connectivity* task (FC: 7, LC: 4, NC: 4), while LC ratings were higher for *indirect accessibility* (FC: 4, LC: 6, NC: 4) and *overview* (FC: 4, LC: 5, NC: 4). The median rating for LC was lower than others for *paths* (FC: 4, LC: 4, NC: 3), while FC was rated lower for *adjacency* (FC: 1, LC: 2, NC: 2).

In terms of traditional visualization values [Sta14], the NC and LC participants gave high marks for *trust, understanding,* and *availability* of information. FC participants rated *timeliness* mildly negative. For the emotional values, we observe generally low affection, the least from the NC group. Here, LC participants rated *intellectual engagement* and *creativity* strongly positive, while *physical engagement* was ranked strongly positive by FC participants.

Qualitative Results. Figure 4.6 shows a comparison of **insights** recorded on-site with insights recorded in the follow-up questionnaire. We classified a total of 274 comments as insights, out of which 149 were recorded on-site and 125 in the follow-up. On-site, we see that FC participants had most insights (FC: 57,LC: 45, NC: 47), and the most deep insights (FC: 17,LC: 13, NC: 6).

On-site, FC participants most notably identified more insights on *adjacency* ("*Zoey and Odin seem to have a lot of friends*"), and *connectivity* ("*seems like Kaia and John are connections between the two groups*"). *Cluster-identification* ("*easily splits in two groups*"), *density* ("*one side is strongly connected*"), and *outlier* identification ("*some animals have only one contact*") were similarly represented in the collected insights across the groups, with FC participants reporting slightly fewer for each category. 59 comments were related to participants' **task load** during the benchmark tasks (FC:25, LC:21, NC:13). 35 comments refer to *mental effort*, while 24 of them point to a increase (FC:11, LC:11, NC:2) and 11 to alleviation (FC:5, LC:5, NC:1). Increased *physical effort* was only reported by FC and LC participants (FC: 4, LC: 1). Six comments refer to *frustration* (FC:5, LC:1). Overall, *effort* was mentioned 8 times—equally often in positive as in negative connotation (FC:2, LC:2, NC:4). Four participants remarked negatively on their performance (FC:3, LC:1). While five FC and LC participants reported increased temporal effort (FC:2, LC:3), five NC participants reported positive experiences in the context. 84 of our transcribed utterances refer to the general **usability** of **NODKANT**.

Layout clarity was the subject of 37 comments in total. Two utterances for each group refer to it positively ("*topology is quite clear*"). The 31 negative comments occurred predominantly among FC and LC participants ("*the messy structure made it a bit harder*") (FC:17, LC:11, NC:3). 17 utterances were classified as expressing participants' *uncertainty* ("*not sure if there was a single shortest path*"). Most of these originated from LC participants (FC:5, LC:11, NC:1). We coded 13 comments as referring to participants recalling details about the graph from *memory* ("*knew the answer from the general overview, just checked if it was right*") during the tasks (FC:5, LC:5, NC:3). Seven statements related the experience during the study to *learning* ("*it is very nice recap for graphs*") (LC:4, NC:3). *Intuitivity* in solving the tasks was predominantly reported by NC participants (FC:5, LC:5, NC:3). We received two comments praising the clarity of our *instructions* by LC participants. A total of 34 comments refer to **engagement** (FC:15, LC:14, NC:5). Of the 14 *affectionate* statements("*playful experience, it was fun!*"), comparably few refer to NC (FC:6, LC:5, NC:3). Similarly, *intellectual engagement* was predominantly reported by FC and LC (FC:5, LC:5, NC:2). Finally, only FC and LC participants reported *physical engagement*. Six participants gave **feedback** in the form of *visual improvement* suggestions, such as coloring nodes and edges.

Summary RQ2: Comprehension. Bearing in mind the stark differences in visual appearance between the different representations, as shown in Figure 4.4, the results above are surprising. Firstly, the **impact of the construction method on task time and accuracy is negligible**. FC participants had to work with their—much messier—graphs, often impacting tasks that are supported by layout clarity, increasing mental effort, and uncertainty , ultimately lessening trust. However, FC **participants reported more and deeper insights** on the first inspection and were physically engaged by construction.

4.4.3 RQ3: Interaction

Analyzing our **recordings** of the on-site study reveals that NC participants never touched the graph and only performed *pointing* interactions. Comparing FC and LC recordings shows that

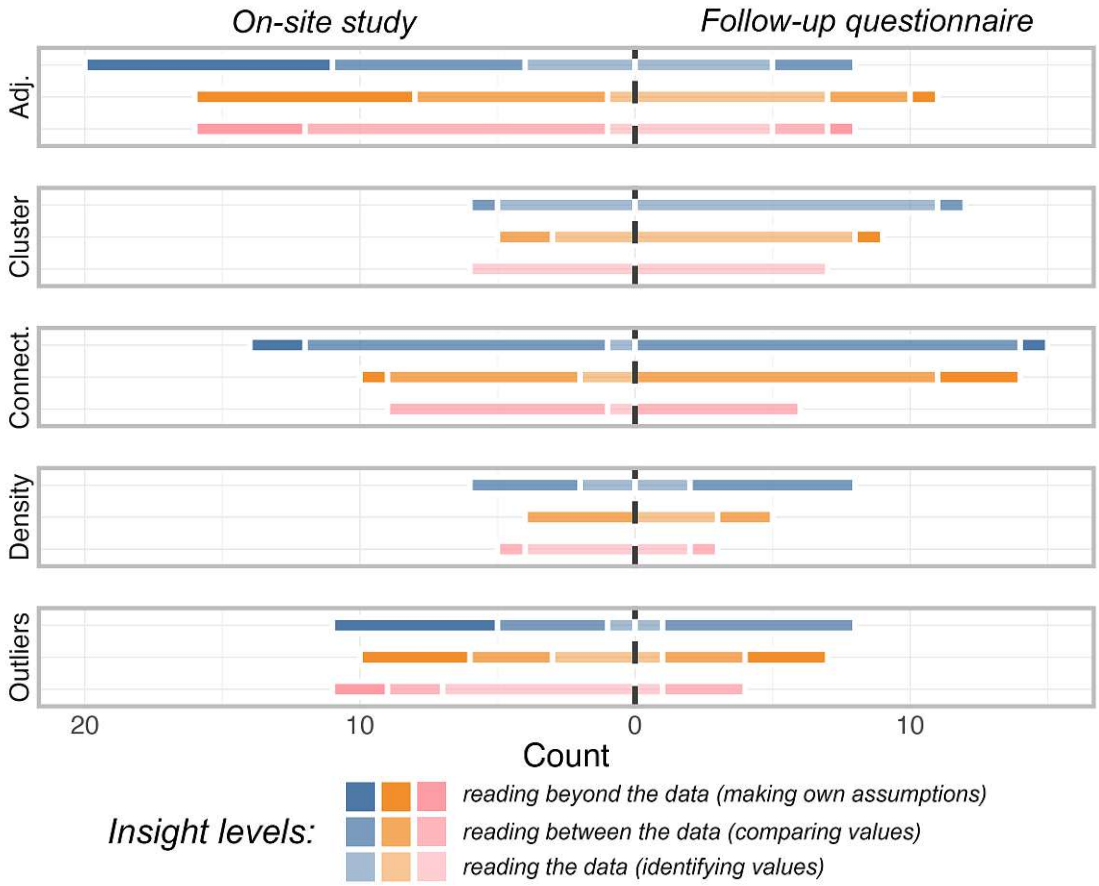

Figure 4.6: Insight codes distribution per group , blue plots relate to FC, orange to LC and pink to NC. Center to left: On-site study insights, with FC participants reporting more insights than LC and NC. Center to right: Follow-up insights, where NC participants report notably fewer insights. Lightness indicates insight level.

FC participants had more *touch*-based contact (FC: 4, LC: 1) and *moved* parts around on the canvas more often (FC: 4, LC: 2). 63 comments related to **physical properties** of **NODKANT**, the least of which come from participants in the NC group (FC:26, LC:29, NC:8). The only physical property referenced by NC was the *height* of the spool stacks, which some participants recognized as an indicator for node degree ("*its nice because the node height indicates degree.*") (FC:12, LC:9, NC:8). FC and NC participants additionally reported problems with *string tangling* (FC:2, LC:2). *String tension* (FC:11, LC:14) and insufficient *string length* (FC:1, LC:4) most often caused problems during the LC condition. Our participants refer to **interaction** techniques to solve benchmark tasks on 86 separate occasions. Most of these comments refer to *visual techniques* ("*went clockwise from the starting node and ticked them off*"), most of which originate from NC participants (FC:8, LC:21, NC:44). Only FC and LC participants report using *physical techniques* to solve the benchmark tasks (FC:9, LC:4). We show the interaction techniques our participant used in Figure 4.7. Four participants moved nodes individually to determine

connections to neighbors ("*I wiggled the nodes and checked if neighbors wiggle*," Figure 4.7 (a)). Three participants refer to manipulating edges ("*tugged on the strings to see the connections better*," Figure 4.7 (b)). Two reported manipulating two nodes at once ("*if you pull two edges apart and see the lines that get straight you can check if they connect to a common neighbor*," Figure 4.7 (c)). Two participants reported comparing node heights by rearranging the graph ("*counted like coins by pushing the stacks close to each other*," Figure 4.7 (d)).

Summary RQ3 : Interaction. Despite explicitly encouraging all participants to physically interact with the representations, we observe that **physical interaction after construction occurred only in FC and LC**. Additionally, despite constructing the representations, **LC participants rarely interacted physically with the graph**. Unsurprisingly, NC participants who did not construct a graph never physically interacted with it. We conclude that **free construction motivates people to physically engage with their representation** more.

4.4.4 RQ4: Memorability

Memorability was investigated in the follow-up study conducted online 10–14 days after the on-site part of the experiment.

Quantitative Results. Average accuracy for **closed questions** was 73.01%, with LC participants' accuracy being slightly lower than others (FC: 75.31, LC: 66.06, NC: 77.91). While most participants across all groups remembered what the *entities* in the graph represented, most NC participants answered only partially correctly on the question on the defined *relation*. Most participants at least remembered one of the given *contexts*. On average, LC participants had lower accuracy in remembering *node count* (FC: 89.58, LC: 70.14, NC: 92.19) and *edge count* (FC: 65.82, LC: 44.11, NC: 77.78), while all three groups did comparably well in deriving the average node degree (FC: 84.72, LC: 83.33, NC: 86.11).

Qualitative Results. Most notably, NC participants reported the least **insights** (Figure 4.6) overall (FC: 51,LC: 46, NC: 28). In general, the insights recorded in the follow-up were classified

(a) (b) (c) (d)

Figure 4.7: Different interactions with **NODKANT**. (a) Wiggling a single node reveals connections. (b) Tugging on nodes reveals common neighbors. (c) Pulling an edge shows connected nodes. (d) Pushing nodes together allows direct comparison of degree.

lower on average, with **LC** participants remembering the most deep insights (**FC**: 1,**LC**: 7, **NC**: 1). **FC** participants retained more insights on *clusters*, *connectivity*, *density* and *outliers*, while **LC** participants recounted more insights on *adjacency*.

Summary RQ4: Memorability. Here, quantitative results show that most participants recalled questions that were asked in the on-site part and could even derive the average node degree relatively accurately, which was not asked before. This attests well to the general memorability of data physicalization. More interestingly, the distribution of reported insights in our follow-up indicates that **constructing a physicalization makes it more memorable**. While we acknowledge that time spent with the representation may be a confounding factor for this, we also observe deeper insights retained in **LC**, which points towards a **positive influence of readability on memorization**.

4.5 Discussion and Conclusion

Comparison with Related Work. Van Ham and Rogowitz [vHR08] found that users tend to enclose clusters in hulls when arranging graphs. We show that such patterns occur similarly in physical settings. Contrary to Purchase et al. [PPP12], our results do not indicate a preference of users for grid-based layouts in the physical space. In addition, our results show that a constructive approach supports the **generation of early and deep insights on the presented data**. We pose constructive approaches like **NODKANT** could support analytical reasoning (Figure 4.8a). Moreover, we show that our toolkit supports physical interaction with our representation and spatial perception of the embodied data. While it does not support the construction of “real" three-dimensional graphs like Bae et al.’s [BFY+24] or McGuffin et al.’s [MSF24] approaches, it uses its 2.5-dimensional properties to convey certain aspects of the graph intuitively. Moreover, as opposed to static contemporary approaches, such as the one presented by Drogemuller et al. [DCW+21], **our technique invites “analog" manipulation** that can be used to navigate network data. While Dwyer et al. [DLF+09] show various interaction patterns of users creating graph layouts on touch screens, **NODKANT** affords unique physical interaction even after construction. While they observed that user-generated layouts could influence performance, our participants were able to compensate for the messiness of their self-generated layouts through familiarity. In the past, physical data representations have been shown to be more memorable than virtual representations [SSB15, HMC+20]. Recently, Pahr et al. [PEW+24] demonstrated that interactivity can enhance the perception of an active physical representation. In contrast, we show that even **interaction during construction makes physicalization more effective and memorable**, presenting a unique opportunity for physical data representations. Future work in data physicalization could make use of constructive metaphors, allowing users to be more involved and creating useful insights into the presented data.

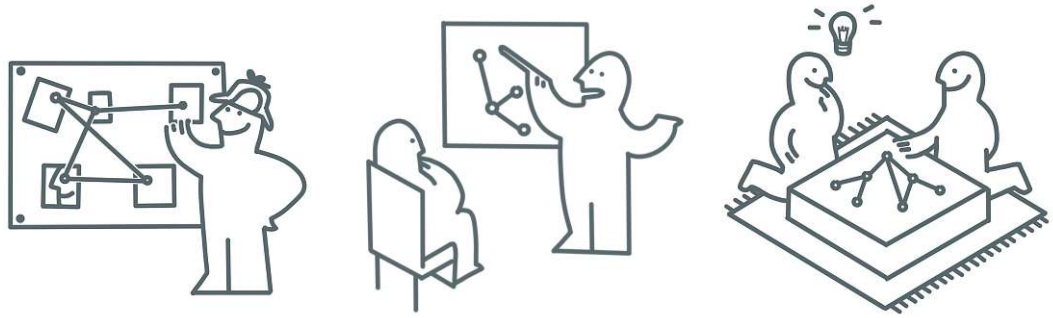

(a) (b) (c)

Figure 4.8: Schematic depictions of applications for constructive network physicalization.

Limitations and Future Work. Our study participants were able to build a small network and completed a series of benchmark tasks with **NODKANT**. We acknowledge that participants often had problems keeping the edges straight, leading to tangling and confusion. While we show that the downsides of this are compensated by physical interactions and memorability, we acknowledge that **users need to be supported during the construction process**. For instance, designing self-retracting edge spools could mitigate the cognitive and physical overhead of construction and allow users to engage more with the data—as opposed to graph aesthetics. We also acknowledge that due to time constraints, only a single network dataset was investigated. As such, investigating how network size and density influence the interaction with **NODKANT** is left for future work. Also, our study population is comprised mainly of young, educated male students. This allowed us to rely on their experience in navigating data visualization. An interesting avenue for future work would be to investigate how the user's involvement while constructing a representation influences their visual literacy, for example in **educational** settings (Figure 4.8b).We also performed isolated experiments with one user at a time, thus limiting the **social aspect** of engagement [WSK+19]. The **NODKANT** toolkit itself does not rely on experience with visualization and only requires a few cheap and easily accessible parts. This could support workshops where groups of people **collaboratively** construct a representation of community-relevant data (Figure 4.8c).

Concluding Remarks. We provide users with a simple, effective, and dynamic toolkit for constructing network physicalization. Huron et al. [HJC14] propose that this allows a user to make use of the "visual mapping" stage of the visualization pipeline [JD13], letting users select an abstract form (i.e. type of chart) for their visualization. By limiting the toolkit to the construction of node-link diagrams, we allow users to engage in the "presentation–mapping", i.e. letting users decide their own layout. **NODKANT** demonstrates that physical interaction with data can enhance the sense-making process through construction by confronting users with data during construction, leading to deeper and more memorable insights and creative interaction with the representation.

CHAPTER 5

Squishicalization: Exploring Elastic Volume Physicalization

This chapter is based on the following publication:

> Pahr, D., Piovarči, M., Wu, H.-Y., Raidou, R.G. (2024). **Squishicalization: Exploring Elastic Volume Physicalization.** In *IEEE Transactions on Visualization and Computer Graphics, vol. 31, no. 9, pp. 6437-6450.* DOI: 10.1109/TVCG.2024.3516481

Various levels of interactivity are seen in data physicalization—from full-body experiences to surface-level tactile interactions. Incorporating data directly into material characteristics like elasticity remains largely unexplored, even though fabrication pipelines that leverage direct interactivity (also in the form of elasticity) exist. However, encoding volumetric data in the internal physical properties of objects has not yet been explored, and we propose that microstructure-based fabrication approaches offer a promising method in this direction. Here, incorporating the tactile sense into the representation is a further step towards integrating a user's actions into data physicalizations. Encoding data into the tactile behavior of an object leads to direct sensory feedback, but the visual sense adds no additional information. (Figure 1.2c).

5.1 Introduction

Data physicalization involves encoding information into physical properties of artifacts [JD13]. Although static physicalizations have long been used for data exploration [DMSJ23], recent approaches have proposed interactive data physicalizations [BZW+22]. As data physicalization goes beyond visually inspecting physical representations and towards multisensory experiences [HH16, HMC+20], the integration of tactile feedback could hold significant promises for tangible data exploration. Tactile feedback offers a means for enhancing interactivity and introduces new opportunities for engagement, allowing users to interact directly with their

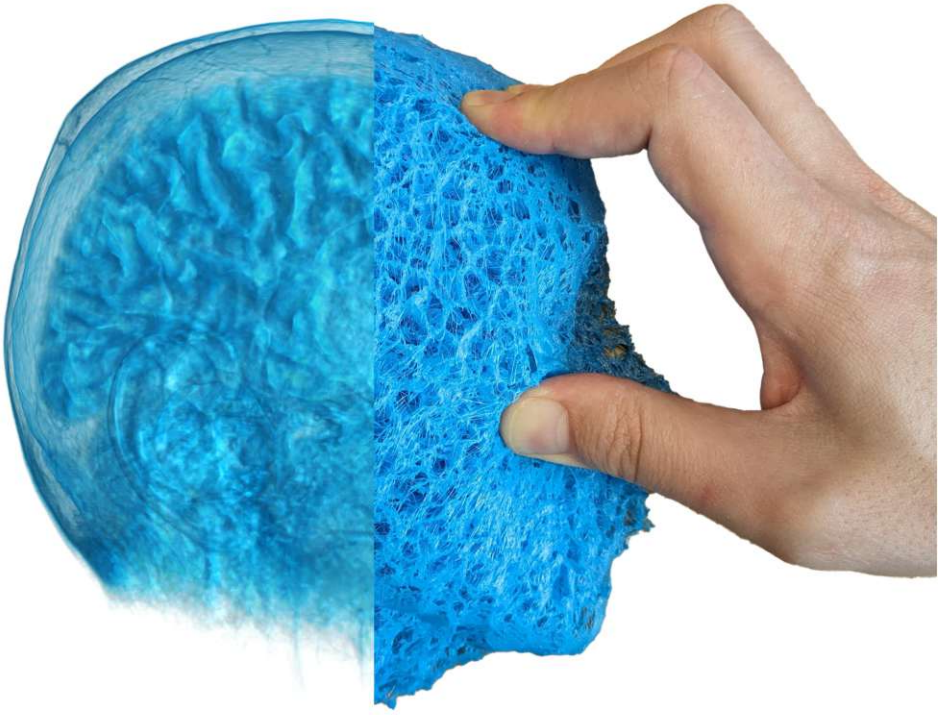

Figure 5.1: A demonstration of the result (right) of the *Squishicalization* pipeline for a volumetric dataset of a human head (left), acquired by Magnetic Resonance Imaging (MRI). Local densities in the squishicalization correspond to scalar values within the volumetric dataset.

data [BCB07]. In data physicalization, tactile feedback is typically realized by making use of the surface properties of objects. Yet, when considering volumetric data, existing research on their physical representation has primarily focused on developing interactive and engaging sculptures that purely use the visual channel [SWR20, RGW20, PWR21].

Inspired by volume rendering, our work seeks to bridge a gap in the literature by introducing *Squishicalization*, a method for translating scalar information in volumetric data into tactile properties of physical artifacts. Similar to how volume rendering techniques map scalar data values to visual properties through transfer functions (TFs), our pipeline extends this concept to density transfer functions (DTFs). These map scalar values of volumetric data onto local density, which subsequently determine the elasticity or "squishiness" of the resulting physical artifacts. This results in squishicalizations, i.e., physicalizations that represent scalar distributions within the data by leveraging physical properties related to elasticity. By adjusting the DTFs, designers control the tactile properties of the physical artifacts to reflect the density distribution of materials within the volumetric dataset. To achieve this, we employ the concept of Voronoi tessellated microstructures [MHSL18] from digital fabrication within our squishicalizations, which are later

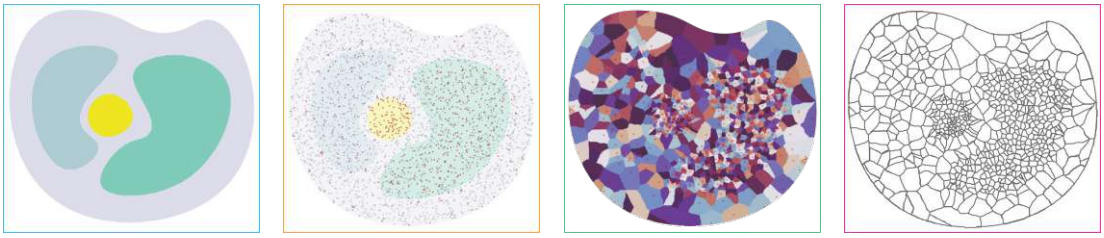

(a) (b) (c) (d)

Figure 5.2: The steps of our *Squishicalization* pipeline. (a) Scalar volumetric data is read from a file or created by a mathematical function. (b) A subsampling (gray: random sample, red: selected seeds) based on the scalar distribution yields seed points to be used for the (c) Voronoi tessellation of the volume. (d) The regions corresponding to the tesselation borders are extracted using iso-contouring.

fabricated with elastic materials using 3D printing. While we show examples created with fused filament fabrication (FFF), any 3D printing method that processes elastic material is equally suited.

Unlike previous works, our approach introduces direct interactivity by allowing users to engage with the data through tactile exploration of the material's elasticity, offering a novel way to explore volumetric data beyond surface structures. With *Squishicalization*, we aim to answer the following **research questions**: (**Q1**) How can we use consumer-grade 3D printing to create tangible artifacts that map volumetric data onto physical properties, specifically elasticity? (**Q2**) How can we control the elasticity of our tangible artifacts to ensure an effective correlation to the scalar values of the volumetric data? (**Q3**) Which practical applications could benefit from elastic volume physicalization?

We design an interactive elastic volume physicalization pipeline, *Squishicalization*, to create tangible sculptures from scalar fields in volumetric data. The individual steps of the pipeline can be likened to those of the volume rendering pipeline. We subsequently evaluate our pipeline's computational, mechanical, and perceptual capabilities. From this evaluation, we distill recommendations for the design of our squishicalizations. We also conducted interviews with expert users from three fields—extended reality (XR), materials science, and medical education—to gather insights into the applicability of our proposed *Squishicalization*.

5.2 Methodology of *Squishicalization*

We present *Squishicalization*, a pipeline for the computer-aided generation of tangible physicalizations that map scalar information from volumetric data onto their internal elasticity. The goal is to directly stimulate the user's tactile sense as a means of direct interactivity [BZW+22].

We refer to **elasticity** as a material's ability to deform by an applied force and subsequent return to its original shape and size. Here, more elastic means that an object or a region in an object deforms more than a less elastic object when the same force is applied. **Stiffness** refers to an object's

resistance to (elastic) deformation. Our volume physicalization pipeline comprises four steps (Figure 5.2) and draws upon foundational concepts of the volume rendering pipeline [Max95], which creates visual representations of volumetric data using properties, such as color and opacity, within well-defined transfer functions (TFs) [LKG+16]. Comparably, our approach aims to generate physical objects from volumetric data that physically encode scalar information onto their physical characteristics—specifically, their elasticity.

We adapt the TF concept to reflect the scalar distribution within the volume onto the distribution of elasticity inside the physicalization. Instead of manipulating color and opacity to enhance structure visibility, we manipulate the "squishiness" based on the scalar values to enhance tangibility. This allows us to create physical objects where the distribution of elasticity maps the distribution of scalar values in the data.

Physicalizing **volumetric data**, as those depicted in Figure 5.2a, is challenging due to the complexity of representing their inherent 3D scalar fields. Slicing and filtering [RGW20, PWR21] have been commonly employed in literature to enhance structure readability. We propose to map the scalar values of a volumetric data set to local elasticity. In a 3D medical imaging dataset, for example, we can represent denser tissues like bones with lower elasticity (i.e., denser means stiffer). This is done by our density transfer functions DTFs, which are comparable to the concept of TFs in conventional volume rendering, and influence sampling densities. Here, a **sampling step** generates a number of seeds per region, reflecting locally the density of the underlying structures (Figure 5.2b). Following up on the medical imaging example above, bones will be sampled with a higher number of seed points than soft tissues. We opt for weighted sampling, to mitigate clustering of the samples [Yuk15]. Subsequently, using **Voronoi tessellation** we divide the volume data into labeled regions using the seed points resulting from the sampling step as the region centers (Figure 5.2c). The Voronoi regions will produce a sponge-like structure, such as the one shown in Figure 5.2d. At this point, dense areas are tessellated into a high number of small regions, whereas sparse areas result in fewer regions of larger volumes [MHSL18]. In the final step, we produce the squishicalizations by **FFF 3D printing**, chosen for its accessibility and cost efficiency [GYM+17]. FFF printing allows the use of elastic materials, which, combined with sponge-like microstructures, can produce soft artifacts. In the resulting squishicalization, denser microstructures can be physically compressed to a lesser degree. This means that our soft artifacts directly encode the scalar values of volumetric data in the form of local elasticity. We elaborate on the steps of our pipeline in the remainder of this section. Several printed examples and their respective elasticities are demonstrated in Figures 5.1 and 5.5.

5.2.1 Data Acquisition and Basic Processing

> We consider a volumetric dataset $V : \mathbf{R}^3 \rightarrow \mathbf{R}$ consisting of scalar values in a regularly spaced three-dimensional grid.

Such data can be obtained in various ways, e.g., through imaging modalities such as computed tomography (CT) or magnetic resonance imaging (MRI). In these cases, the scalar values correspond to measurements of physical phenomena, such as the absorption of X-rays in materials or the response of molecules to magnetic polarization. These scans vary in voxel size, from

around 1*mm* in full or partial body scans to much smaller in tissue scans, where individual cells can be viewed. Although prevalent in medicine, **3D image data** are also used in other fields, e.g., in materials science. Alternatively, scalar volumetric data can be synthetically generated by **mathematical functions** in the form $f(x,y,z) = s$, which can be sampled over discrete points in a given interval. The result can be visualized in the 3D space, where each voxel encodes the function value at a given point. The resulting volumetric data are commonly viewed on screens as axial slices, using volume rendering techniques [Max95], or by extracting one or multiple isosurfaces [LC98].

> A typical step in volume rendering is **thresholding**. It can be used to segment volumetric data by eliminating unwanted scalar values within user-determined ranges (e.g., under a certain value). This way, artifacts from data acquisition can be removed or the background of an image can be ignored, reducing the amount of data to be processed. By extracting an **isosurface** from the scalar field, we can additionally create a 3D mesh of a structure. Such meshes can be combined in the printing process to add solid parts, achieving the upper threshold of stiffness for a given material, bypassing the need to create a microstructure.

5.2.2 Weighted Volume Sampling

Our approach aims to transform our volumetric data into regions of different elasticity, depending on their scalar values. Previous work by Lu et al. [LSZ+14] deals with the fabrication of solid models based on a force distribution, while Martinez et al. [MHSL18] investigate the generation of elastic models based on given point distributions. In contrast, our method introduces an customizable approach for generating point distributions directly from diverse volume data, based on the following simple concept: a continuous range of scalar values (e.g., each of the four colored regions in Figure 5.2a) should correspond to a specific stiffness within the final physical object. Therefore, we adapt the approach of Martinez et al. to create point distributions based on the underlying scalar values of the volumetric data. We design our approach **independently of prior segmentation**, and directly rely on the scalar values.

The tesselation is discussed in the upcoming section, but its basis requires a set of points, which can be produced through sampling that follows these three requirements (Figure 5.2b): First, **regions with similar scalar values should have similar elasticity**. Thus, their underlying tesselations or cells should be sized approximately equally to ensure similar elastic properties [MHSL18]. Limiting the number of points used to represent a given volume, or adjusting the sampling probability based on scalar values does not guarantee this. For instance, Lu et al. [LSZ+14] use an error diffusion approach to create their sample distribution, which also does not guarantee regularly spaced sampling points.Random sampling can introduce undesirably clustered points in regions where the scalar field distribution is stable, introducing tangible artifacts instead of reflecting the scalar field accurately.

Second, **point density of a region should relate to its scalar values**. Ideally, the relationship between scalars and sampling density should be customizable to allow flexibility in manufacturing. Here, we call back to the application of TFs in volume rendering, and we employ a similar concept to relate scalar values to elasticity through sampling density (density TF or DTF, further described

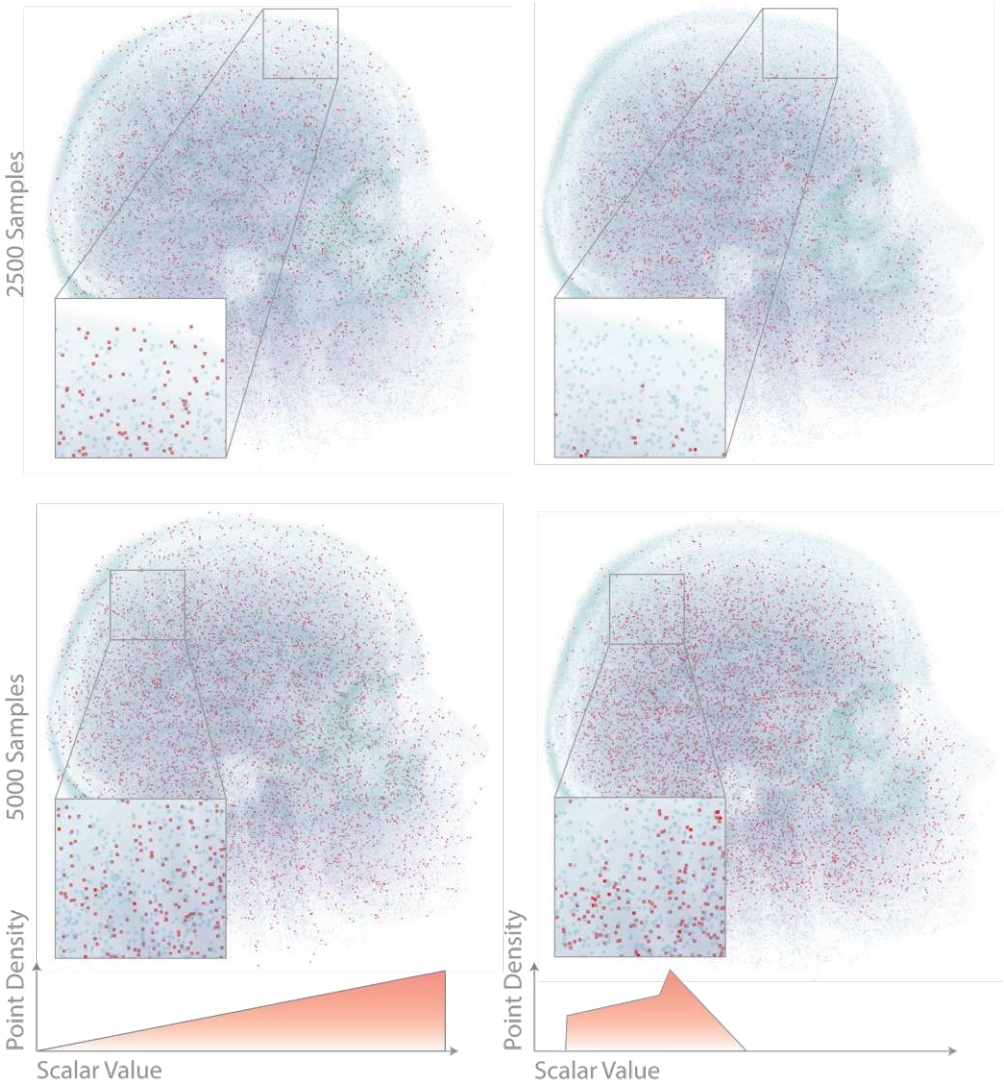

Figure 5.3: Weighted sampling of a volumetric dataset for different DTFs. The volume rendering of a head CT dataset is overlayed with gray points (sampling input) and red points (sampling output). In the left column, a simple linear function yields fewer selected samples in softer areas like the brain and more selected samples in harder regions like the skull. In the right column, we show a custom DTF that leads to more selected sample points from cerebral structures, and fewer chosen samples from regions of the skull.

below). Third, we need to **limit computational complexity** as much as possible. Tessellation is a computationally expensive process, whose performance depends on the selected number of sample points. In our case, this is further complicated by the dimensionality of our data.

A method that fulfills all of these requirements is **weighted Poisson-disk sampling** [Bri07]. In this method, a minimum distance between two sample points is defined to create an equally

spaced point set. This distance can be a function of a local weight, and in our case, this relates to the scalar values of the data. However, the performance of Poisson-disk sampling algorithms is impacted by the dimensionality of the sampling space, as well as the number of candidate points. Yuksel [Yuk15] proposes a sample elimination approach to create sets with Poisson-disk characteristics from a given point set. This method enables us to create sample sets with equal spacing of points, while the minimal distance between points can be a function of the local scalars. A random sample of points of a sufficient size inside the original volume can serve as the input for the sample reduction. The size of the desired sample is also a parameter of our algorithm.

> To obtain a set of sample points we use a sample reduction technique that **iteratively eliminates points in a given set until a desired sample size is reached.** We follow Yuksel's approach [Yuk15] for the sample reduction, using the recommended parameter values $\alpha = 8$, $\beta = 0.65$, and $\gamma = 1.5$. A higher α increases the influence of point distance on the weight function, while β and γ values are used to lessen the influence of very close points compared to far-away ones. The approach also allows for the use of a weight function that considers local scalar values in addition to point distances.

For the customization of the weighted sampling, we propose the use of **density transfer functions (DTFs)**, an analogous concept to color and/or opacity TFs in volume rendering. We define a DTF as a function $f(s) = d_{max}(s)$, where s is a scalar value in a volume dataset and d_{max} is the maximum radius for the given scalar value. Conventional TFs make use of the concept of **piecewise linear TFs** [LKG+16]. As we do not know the precise relation between sample density and elasticity and users may prefer to emphasize certain structures over others (Figure 5.3), piecewise linear functions are also suited for DTFs.

> **We define a DTF** by supplying a set of control points $\mathcal{C} = C_i$, consisting of value pairs $C_i = (s_i, d_i)$, where d_i is the maximum radius parameter for a given scalar value s_i. The maximum radius for a scalar value at a given s is calculated as the linear interpolation between d_k and d_{k+1}, where $s_k < s < s_{k+1}$.

Figure 5.2b schematically depicts the weighted sampling concept for four regions with different scalar values, Figure 5.3 shows a concrete example of the sampling process for two different DTFs.

5.2.3 Polyhedral Voronoi Structure Generation

After obtaining the set of sample points $\mathcal{C}$, we create a tessellation from the original volumetric dataset (Figure 5.2c). From this, we extract the desired sponge-like microstructure. For the tessellation, we pose the following requirements: First, our goal is to create **printable sculptures** that can be produced using FFF, which allows the printing of elastic objects using specific materials. Such 3D printers create objects by adding layer-by-layer to the printing area. Yet, large overhanging areas have to be supported by creating support structures underneath, which would influence the mechanical properties of the print. Also, the microstructures we aim to create do not allow the extraction of support structures after print. We, therefore, require that printing the

microstructures does not require support. Second, squishicalizations should exhibit **isotropic elasticity**. This means that the tessellation should not influence the mechanical properties in a non-uniform way (e.g., along a specific axis as in the case of Martinez et al. [MHSL18]). Otherwise, this would result in inconsistent behavior when exploring artifacts encoding scalar values into local densities, where the direction of compression could significantly impact tactility. Third, **shapes** inherent to the volumetric data should be preserved as much as possible. This could mean that the shape of the anatomical area depicted in medical images can be reproduced with printing. Additionally, limiting the volume to "visible" areas (e.g., by suppressing the background with a threshold, as discussed in Section 5.2.1) can increase the speed of the operation performed significantly, since "invisible" areas do not have to be tessellated.

> We use a **Voronoi tessellation**-based approach to create elastic objects with customizable density. In a **Voronoi tesselation** of a volume, given a set of sample points, each voxel contains the label of its assigned sample point. A **distance function** determines the proximity of all samples to a single voxel. The labels are chosen by finding the closest sample point.

Martinez et al.'s [MHSL18] approach uses **polyhedral distance functions**, based on parametrizable cones. They argue that, when using a Euclidean distance function, the borders of the resulting areas could have arbitrary angles, causing overhangs that cannot be printed without support structures. In contrast, their customizable cone-shaped distance function can limit border angles resulting in the tessellation to exclude steep drops. We fulfill our first requirement, by using a similar polyhedral distance function for the tessellation.

> **A polyhedral distance function** is defined as follows: To calculate the distance between points P_1 and P_2, we construct a convex polyhedron that contains $P1$, defined by a set of planes $\mathcal{B}$ traversing the vertices of the polygon, defining its faces. We first calculate the intersections with every plane on a connecting line between P_1 and P_2. The closest intersection point I_{min} to P_1 indicates where the connecting line intersects with the polygon. The polyhedral distance between P_1 and P_2 is the ratio between the Euclidean distance between the points and the distance between P_1 and I_{min} so that $d_{\mathcal{B}}(P_1,P_2) = \frac{|P_2-P_1|}{|I_min-P_1|}$. This corresponds to a **scaling factor** applied to the given polygon around P_1 so that it just touches P_2, without rotating the polygon. In the case where the line between two candidate seed points is parallel to a face of the polygon, i.e., the two computed distances would be equal, we opt for the seed point with the lower index.

From our second requirement, we do not desire customizable anisotropy, i.e., we want to avoid adjusting directional variations in the properties of the model. Thus, we use the bitruncated cubic honeycomb polyhedron. This way, the **distance function guarantees isometry** along every axis. As a honeycomb shape, this polyhedron can be arranged to create a tesselation of 3D space using only translation. It also contains only steep angles and horizontal planes, small enough to bridge small gaps to reduce the need for support—in accordance with the first requirement. To fulfill the third requirement, we limit the seed points for the tessellation by defining a background scalar value, which will be suppressed and excluded from both subsampling and tessellation. For example, in Figure 5.2 the white region around the volume is not considered. This limits

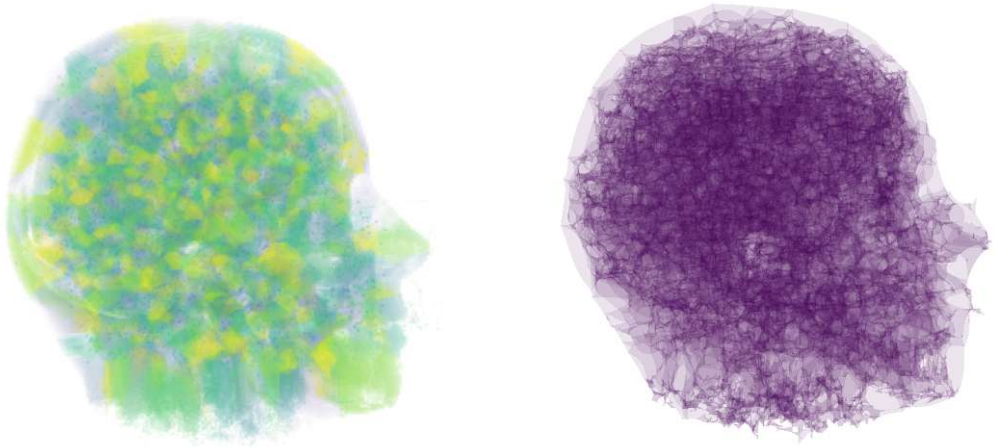

(a) (b)

Figure 5.4: (a) Labelled regions and (b) extracted border mesh, created by Voronoi tessellation with a bi-truncated honeycomb distance function. Sample point set taken from Figure 5.3, bottom right.

performance impacts and produces a volume that **preserves the shape** of a region of interest. Figure 5.2c schematically depicts this process, while we show a labeled tessellation performed with this method in Figure 5.4a. Figure 5.9 shows examples with different regular point densities.

5.2.4 Meshing

To create a printable mesh, we need to obtain the region borders from our tessellation result (Figure 5.2d). Martinez et al. [MHSL18] investigate a layerwise path extraction. These methods create a large performance bottleneck. As the Voronoi regions in the mesh are already labeled with integer values, we opt for a labeled iso-surface extraction approach [Fri22] instead:

> To extract the region border we use an implementation of Frisken's surface nets algorithm [Fri22] in PyVista [SK19]. It produces a single mesh of region borders for a labeled volume. Because we also obtain a label for the background voxel value in the first step, the object's outer shape is preserved. In addition to the extraction of borders, this approach also produces smoother edges and more equally spaced vertices, removing undesirable staircase artifacts and simplifying the generated geometry.

Figure 5.2d schematically depicts the mesh extraction process, while we also show a 3D result in Figure 5.4b.

5.2.5 Slicing and Printing

In the last step, we process the resulting mesh using readily available slicing software. After preparation, the squishicalization can be printed. Because a multitude of 3D printers and a wide

range of materials exist, we require this step to be **independent of printing hardware** and as **customizeable** as possible. For accessibility, we also require the **materials to be compatible** with consumer-grade 3D printers. Slicing software serves to transform triangle meshes into toolpaths. These are created by cutting a mesh along the vertical axis and analyzing the geometry on the cutting plane. Many aspects of this process are customizable, such as the selection of printer, wall thickness, and parameters of support structures. Such software is usually designed for closed, non-manifold meshes, as a slicer needs to determine the inside and outside of a mesh to create infill patterns that are used to limit filament usage while preserving the mechanical stability of a 3D-printed object. However, multi-label segmentation is not designed to produce such well-behaved meshes. *UltiMaker Cura* provides a **surface mode** that processes a mesh independently of composition and extracts any borders in a mesh as a single wall to print. *UltiMaker Cura* supports toolpath creation for a multitude of contemporary 3D printers. The use of slicer software also allows us to **combine different meshes** in a single print. This way, solid objects can be suspended in soft material, increasing the range of elasticity we can achieve. We use **thermoplastic polyurethane (TPU)** as it is commercially available in a wide range of colors, from different manufacturers, and compatible with regular FFF 3D printers. An example of the head dataset is demonstrated in Figure 5.1.

5.2.6 Implementation

We implemented our pipeline in *Python*, with slicing handled externally by *UltiMaker Cura*. The weighted Poisson-disk sample reduction algorithm proposed by Yuksel [Yuk15] was extended to be compatible with the concept of DTFs using *NumPy*. We parallelize the calculation on the GPU using *Numba* to create *CUDA* kernels for the polyhedral voronoi tessellation approach proposed by Martinez et al. [MHSL18]. An implementation of Friskens' [Fri22] surface nets algorithm in *PyVista* serves for surface extraction and mesh smoothing. Our implementation is available at *OSF*.

5.3 *Squishicalization* Results

We now show some results of our physicalization pipeline and test the capabilities of our approach in a computational, a mechanical, and a perceptual evaluation. For more details on the tasks, procedure, results of the study, and prop descriptions, we refer the reader to our supplemental materials at `https://osf.io/35gnv/?view_only=605e5085061f40439a98545f0c447cf3`.

5.3.1 Performance

We tested different parameter combinations and evaluated the impact of individual settings on the performance of *Squishicalization*, using a total of four datasets (Table 5.1). The **generic** dataset consists of $300 \times 300 \times 300$ voxels with randomized scalar values. The **foot** dataset is a CT scan of a human foot and the **head** dataset is an MRI of a woman—both taken from the Open SciVis Datasets. The **fibers** dataset was provided by a domain expert in the later course of an interview, described in Section 5.4. We limit the parameter space exploration of the performance to the parts of the pipeline we implemented (Sections 5.2.1– 5.2.4), as slicing (Section 5.2.5) is

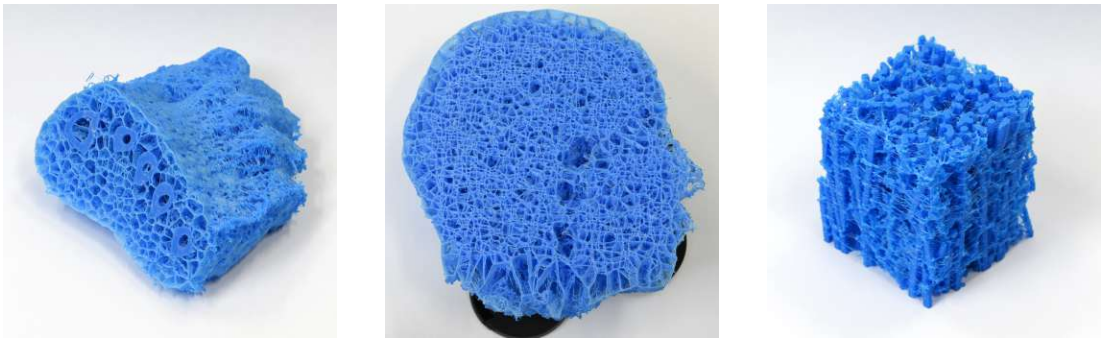

(a) (b) (c)

Figure 5.5: Resulting squishicalizations for different volumetric datasets printed on a Prusa i3 Mk3s using an opaque TPU filament. (a) Foot CT dataset from the Open SciVis Datasets. (b) Head MRI dataset from the Open SciVis Datasets. (c) Industrial CT dataset provided by a materials science expert.

performed in external software. The tests were performed on a system with an AMD Ryzen 9 3900X 12-Core Processor with 3800 Mhz, an NVIDIA GeForce GTX 1080, and 64GB RAM. We show the results in Table 5.1. Besides the FFF 3D printing bottleneck, our benchmarks show that the tesselation is the most time-consuming step, taking between $0:53$ and $8:27$ minutes. The sampling becomes an intense factor for big sample sizes. While processing the head dataset, we created a sample of $20k$ from $60k$ random points to account for the size of the final print. Sampling took $17:09$ minutes, while tessellation of the whole volume took $15:09$. Limiting the tessellation to voxels with values greater than a predefined threshold reduced the time for tessellation to $2:18$. Mesh creation is always negligible (approx. one second).

All prints were done on a Prusa i3 Mk3s, with a wall thickness of $0.2mm$. For solid parts, we used 100% infill density. Regarding **printing times** and **material use**, the foot dataset on a scale

Table 5.1: Performance for different datasets. *Head** marks the use of a scalar threshold, reducing the impact of the tessellation performance. n_r indicates the number of random samples used in the sample elimination and n_s indicates the number of seeds, t indicates measured time ($m:ss$).

Dataset	**Size**	**Sampling**			**Tessel.**	**Meshing**	**Total**
		n_r	n_s	t	t	t	t
Generic	300×300×300	1k	0.3k	<1s	00:53	<1s	0:54
		5k	1.5k	0:01	4:10	0:01	4:12
		10k	3k	0:04	8:27	0:02	8:33
Foot	256×256×256	10k	3k	0:12	1:32	<1s	1:45
Fibers	350×350×400	10k	0.3k	0:13	1:36	<1s	1:51
Head	256×256×70	60k	20k	17:09	15:09	<1s	32:20
Head*	256×256×70	60k	20k	17:11	2:18	0:01	19:31

of $10cm \times 10cm \times 10cm$ took about one day and 5 hours, using $78g$ of filament. The foot was printed combined with a mesh created from an isosurface extracted the scalar values representing solid bones. In the DTF, only soft tissue is considered in the sampling process and bones are initially excluded; and then, reinserted as an isosurface. The outcome is shown in Figure 5.5a. The fibers dataset took about 19 hours to print on a scale of $5cm \times 5cm \times 6cm$, and used $44g$ of filament. Here, we also used an isosurface of the fiber material in the print and a dense sampling of the material in which the fibers are embedded. This allows the fibers to move more freely while still maintaining a coherent structure. The outcome is shown in Figure 5.5c. Finally, the head dataset, at a scale of $23cm \times 20cm \times 9m$, took two days and 7 hours to print and consumed $194g$ of filament. The DTF was designed to emphasize a tactile difference between gray and white matter, while harder tissues were sampled more sparsely. We did not include solid parts in the print. To save time and material, and to facilitate an inside view of our microstructure, we clipped the dataset along the coronal axis and excluded all scalar values under 7% of the maximum, to limit noise and performance impact. The *squishicalization* of the head dataset is shown in Figures 5.1 and 5.5b.

Upon visual inspection, our test prints show that vertical walls are well pronounced. As the structures approach more flat angles, staircase artifacts consisting of disconnected strings become visible. This can be explained by the slicing technique we employed to process the meshes. *UltiMaker's* "Surface Mode" produces a single wall for every intersection of the mesh with the horizontal plane. When the distance between two sequential layers becomes too large, this leads to layers becoming disconnected at places. We further discuss the impact of these artifacts in later sections.

5.3.2 Perceptual Study

We conducted a study with 18 participants to investigate the abilities of our squishicalizations to stimulate tactile perception. The objective of our study was to investigate whether it is feasible for humans to recognize and interpret complex 3D structures by directly interacting with their respective squishicalizations. We are particularly interested in properties such as partitions, gradients, and the presence of heterogeneous regions within 3D synthetic stimuli.

Participants. We recruited our participants from university staff and students without offering compensation. Our participants (male: 16, female: 2) were between 22 and 36 years old (Mean=29, Std=4.05). Two hold a Ph.D., 14 a M.Sc., one a B.Sc., and one has a high school education.

Tasks. Our tasks were designed based on three objectives: *recognizing partitions* within the volume, i.e., part A is less elastic/denser than part B, *elasticity gradients*, i.e., directional changes of elasticity, and *position and size* of heterogeneous regions in objects with variable density. In the **partitions** task, participants had to decide how a cube consisting of two equally sized regions differing in stiffness was partitioned. We handed participants one cube at a time, each cube split along a different plane. The **gradients** task had participants determining the direction of a stiffness gradient running through a cube. Again, we handed them one cube at a time, each created

with a stiffness gradient along a different diagonal. Finally, in the **regions** task, participants had to determine the diameter and position of a spherical region with a higher stiffness than the surrounding volume.

Datasets. 3D scalar data with varying partitions, gradients, and heterogeneities can be easily expressed by 3D arrays, simplifying the creation of the datasets for the study objects. Hence, we created three cubes with 50*mm* side lengths for each task. An additional top and bottom thin layer was added to the cubes to visually obscure the inner microstructures. All samples were printed with 0.3*mm* wall thickness using an opaque TPU filament. For the **partitions** task, we created synthetic cubes with two regions of distinct densities. These were parted in the middle along either one of the coordinate axes (`P1`), one of the face diagonals (`P2`, see Figure 5.6a), or the body diagonal of the cube (`P3`). In the **gradients** tasks, we used synthetic cubes with a linear density gradient. We created three datasets, with increasing values along one of the coordinate axes (`G1`), one of the face diagonals (`G2`), and the body diagonal of the cube (`G3`, see Figure 5.6b). Lastly, for the **regions** tasks, we created synthetic cubes of homogeneous density containing an additional spherical region with a different density than the rest of the cube. We used three different diameters for the regions, 20% (`R1`), 33% (`R2`, see Figure 5.6c), and 50% (`R3`) and placed the regions in randomized positions within the cube.

Procedure. Each participant performed all three tasks with three cubes per task presented in a balanced, pseudorandom order. Every cube was presented to the study participants an equal amount of times in all three steps in each task by using a Latin square scheme. Nine of the participants completed the tasks in order (partitions, gradients, regions), and nine completed them in reverse order to combat the influence of learning effects. We measured the **time** participants spent on each experiment. Additionally, participants were asked to rate perceived task **difficulty** and the **confidence** in their answer on a scale of 1–10, ranging from low to high. For the

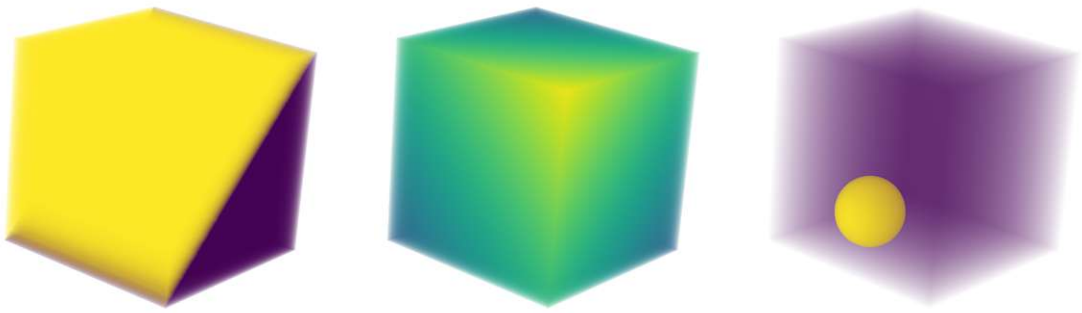

(a) (b) (c)

Figure 5.6: Three different datasets were used in our perceptual study. Different colors represent different scalar values. (a) Dataset for the partition task `P2` (face diagonal partition), (b) Dataset for the gradients task `G3` (gradient along body diagonal), (c) Dataset for the regions task `R2` (medium-sized region)

partitions and the **gradients** task, we gave the participants **single-choice questions** with 7 possible answers, meaning a 14.23% chance of randomly guessing the answer. The possible answers in the partitions task were three axis-parallel partitions, three partitions along the face diagonal, and the body diagonal of the cube. In the gradients task, the possible answers were the three gradients along the coordinate axes, three along the face diagonals, and one along the body diagonal. In the regions task, we asked the participants to indicate the **position** of the center and the **diameter** of the denser region in the cube, measured proportional to the cube side length. We let our participants freely interact with our study props. We also asked participants for **qualitative feedback** after each task and after they completed all three tasks.

Study Results. For the partitions and gradients task, we report the accuracy of the answers of our participants for each dataset. For the regions task, we report the mean deviation of our participants' prediction of the center and diameter of the spherical region from the ground truth in percent. We also show the results of the regions task in Figure 5.7. For all tasks, we report the average time taken, perceived difficulty, and confidence. The results are shown in detail in Table 5.2.

Analyzing the **quantitative** results shows that some participants had trouble recognizing partitions

Table 5.2: Quantitative study results for the partitions and gradients task. Times in $m:ss$, accuracy for positions and gradients in the percentage of total right answers, region positioning error in percent of cube side length, and region size error shown in relation to cube size (both incl. $\pm$ standard deviations and the times, difficulty, and confidence are indicated for both positioning and size estimations together).

Dataset	Accuracy	Time		Difficulty		Confidence	
		avg	std	avg	std	avg	std
P1	50%	1:19	0:42	4.83	2.75	6.61	2.45
P2	17%	1:31	1:03	4.72	2.61	6.83	2.12
P3	22%	1:45	1:04	5.56	2.36	5.94	1.92
P1-3	**29.63%**	**1:32**	**0:57**	**5.04**	**2.55**	**6.46**	**2.17**
G1	39%	1:47	0:49	6.06	2.58	5.78	2.18
G2	11%	1:51	1:03	5.94	1.86	5.50	1.58
G3	22%	1:55	1:03	5.50	1.58	5.67	2.00
G1-3	**24.07%**	**1:51**	**0:58**	**5.98**	**2.11**	**5.58**	**1.87**
R1 pos.	38% $\pm$ 18%						
R1 size	4% $\pm$ 10%	2:22	1:13	6.28	2.63	5.33	2.35
R2 pos.	15% $\pm$ 12%						
R2 size	-4% $\pm$ 17%	1:45	0:29	4.22	2.24	7.17	2.09
R3 pos.	16% $\pm$ 9%						
R3 size	-2% $\pm$ 22%	2:25	1:04	5.56	2.09	6.28	1.99
R1-3		**2:11**	**1:00**	**5.35**	**2.44**	**6.26**	**2.24**

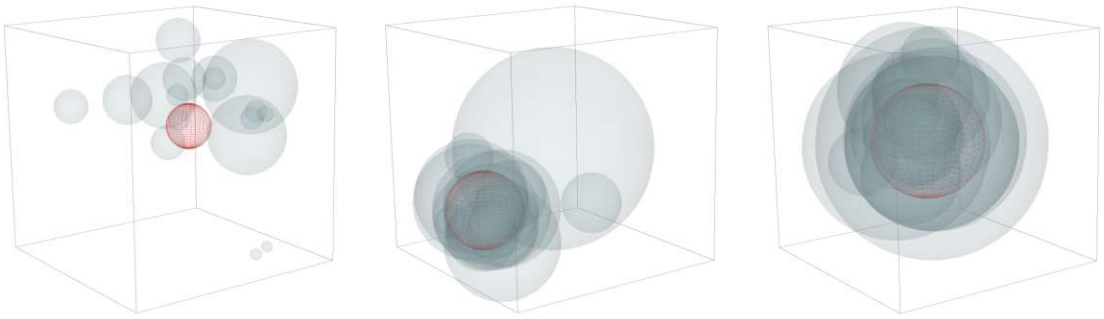

(a) (b) (c)

Figure 5.7: Study results for the regions task. The red sphere in each picture represents the ground truth, while each grey sphere represents the answer of one participant for each of the three cases ((a) `R1`,(b) `R2`,(c) `R3`).

(17%) and gradients (22%) along face diagonals. They were only slightly more accurate in recognizing partitions and gradients along the body diagonal (both 22%). Half of the participants correctly indicated the partition along the coordinate axis (50%) and slightly fewer participants were able to identify the axis parallel gradient (39%). On average, participants rated the difficulty of the gradients task higher than the partitions task (5.98 vs 5.04), whereas indicated confidence was lower (5.58 vs 6.46). Figure 5.7a shows that participants were not able to accurately determine the position of the smallest region (`R1`) in the regions task. However, with an increased size of the region in `R2` and `R3` (Figures 5.7b, 5.7c), the predictions get increasingly accurate. On average, participants slightly underestimated the size of the region in `R2` (-13%) and `R3` (-4%).

The **qualitative** feedback from the participants gives additional insights into why especially the partitions and the gradients task were perceived so difficult and had such low accuracies. Most participants (67%) pointed out at some point during the study that they were irritated by the lower elasticity of the cubes along the z-axis. A number of participants (28%) also remarked that the surface of the cubes was uneven or that the sides, bottom, and top of the cubes were very stiff. Some participants (22%) also directly referenced that the region in `R1` was too small or too hard to find. One additional participant (6%) suspected there was no region of higher density in `R1` and gave an invalid answer. Four participants (22%) directly mentioned not being used to tactile perception before the experiment. A few other participants (11%) also mentioned that they felt fatigued during the experiment, but did not quit the experiment. Finally, the majority of participants (67%) found the experience "*cool*", "*fun*", "*pleasant*", or "*interesting*".

5.3.3 Mechanical Evaluation

The next step in our evaluation is to assess the mechanical performance of the squishicalizations through uniaxial loading tests [MDL16, PLR+16]. These tests aim to understand how the samples respond to compression forces and measure their deformation behavior under such conditions (Figure 5.8). As the testing geometry, we opt for a metallic cylinder with an 80*mm* diameter to ensure consistent force application over the whole sample. The geometry is lowered towards the

sample at a constant velocity of $0.2mm/s$. The sample is compressed until $25mm$ of displacement or until reaching $100N$ of force — whichever happens first. We test and measure each sample along all primary axes to understand how the sample deforms in different directions. This aims to provide insights into the mechanical behavior of the samples, which can inform their suitability for various applications or provide feedback for further design improvements. We show the force–displacement perpendicular to the printing direction observed during the loading tests in Figure 5.9.

Comparing force–displacement curves to evaluate elasticity is highly challenging. To enable quantitative analysis of the measured data we employ the perceptual model for elasticity proposed by Piovarci et al. [PLR+16]. The model proposes a one-dimensional space of elasticity ordering measured samples from softest to stiffest. Furthermore, the perceptual space can be anchored by selecting a reference object. This allows us to express the difference in perceived stiffness between the reference and other objects as so-called Just Noticeable Differences (JNDs). When two samples are 1 or more JND apart they are easy to perceptually distinguish by human observers. Below 1 JND the samples get progressively harder and are generally considered indistinguishable at 0.5 JND. Applying the model to our samples anchored on the softest material reveals that the span between the softest and hardest material covers a total of 287 JNDs. The difference in elasticity in the two directions perpendicular to the printing direction is challenging to distinguish. However, the difference in stiffness along the printing axis reveals a significant increase in stiffness on average by 21 JNDs. This change is caused by both the different behavior along the printing axis and the outer wall of the cube that introduced significant buckling into the samples (see inset). Lastly, to visualize the range of achievable material properties with our system we compare them against the measurements of human tissues gathered by Guimaraes et al. [GGMR20], as depicted in Figure 5.10. We observe that the manufactured samples span a wide range of mechanical properties ranging from almost mucus-like to cartilage-like tissues. We discuss the implications of our mechanical analysis in Section 5.5.

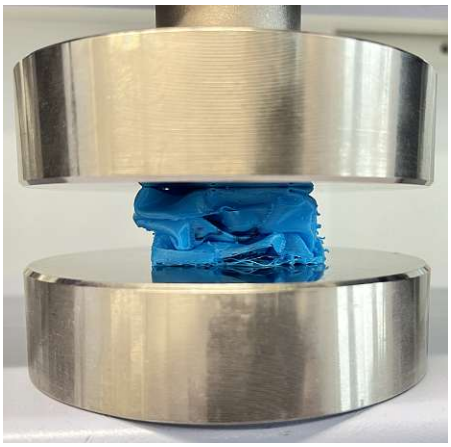

Figure 5.8: Assessing the mechanical performance of a squishicalization through uniaxial loading.

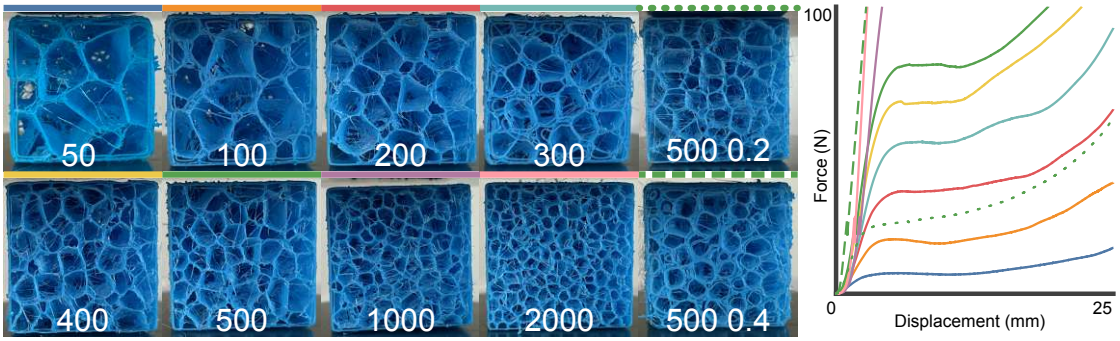

Figure 5.9: Force–displacement measurements for the manufactured squishicalizations. The samples were fabricated with different seed point counts and a fixed wall width of 0.3*mm*. For comparison, we manufactured one sample in 0.2 and 0.4*mm* wall thickness (in the last column).

5.3.4 Follow-Up Study

With the results of our perceptual and mechanical studies in hand, we set out to further support the feasibility of our method. This time we focused on different aspects of tactile perception in addition to improving the fidelity of the props according to our prior findings.

Participants. We again recruited 18 participants, 10 of which had participated in the previous study. The participants were between 24 and 42 years old (Mean=30.2, Std=5.29). Three hold a Ph.D., 13 a M.Sc., and two a B.Sc. 12 of the participants identify as male and 6 as female.

Tasks. The tasks were designed around *sorting objects according to their stiffness* to determine the ability to recognize different levels of elasticity, *recognizing compositions of different complexity* to determine if one can perceive how many different parts of varying elasticity an object is made up of, and *recognizing elasticity variations in objects* to determine the perceptibility of elasticity variations in nested objects. In the **sorting** task, participants had to sort three spheres based on their stiffness. All three spheres were handed out at once and completion time was measured. For the **counting** task, participants were presented with one homogeneous sphere, one consisting of two nested regions, and one consisting of three nested regions (Figure 5.11).

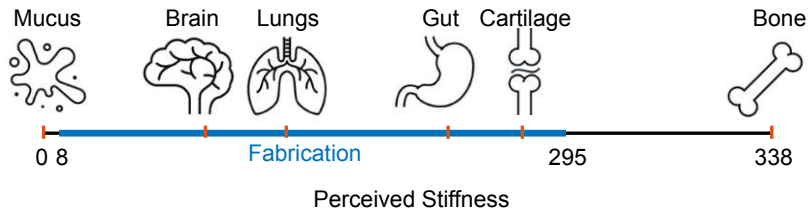

Figure 5.10: Perceived stiffness range covered by real human tissues superimposed with the range of our squishicalizations (in blue). The perceptual space is anchored with mucus as the reference. One unit of perceived stiffness corresponds to one Just Noticeable Difference (JND).

Again, we handed the participants all three spheres at once and asked them to assign each sphere to a composition (1, 2, or 3 regions), with each composition occurring exactly once. In the **assignment** task, we handed the participants spheres, each consisting of three distinct regions, and let them determine their composition. Each region within the sphere had a different level of stiffness ("soft", "hard", and "medium"; similar to Figure 5.11c). The task was divided into two phases, each employing three spheres. In the first phase, we provided the spheres' orientations to assess whether this information aids in determining the sphere composition. In the second phase, participants had to identify the alignment without any orientation cues. Lastly, we handed the participants one of our three prototype squishicalizations at a time and encouraged them to **exploration** freely, sharing their observations.

Datasets. In our first study, the outer layer of the cubes, especially on the top and bottom was perceived as confusing by many participants and the wall thickness used to create the props led to a noticeable difference in elasticity along the print axis. To combat this, we designed spherical props with a 75*mm* diameter, which do not have vertical or horizontal surfaces that could influence perception like cubes do. Additionally, we opted to use a wall thickness of 0.2*mm* as opposed to 0.3*mm* to reduce the anisotropic effects along the print axis. We used three different datasets, composed of one single homogeneous region (Figure 5.11a), two nested regions (Figure 5.11b), or three nested regions (Figure 5.11c) with varying stiffness. We created multiple physical representations of each dataset by varying the compositions in terms of elasticity. For the homogeneous dataset, we created three versions with increasing stiffness. For the two regions, we created both permutations of the dataset: soft outside and hard inside, and vice versa. For the three regions, we created representations of all six permutations of combinations of three different levels of stiffness.

Procedure. Participants completed the tasks in the order sorting – counting – assignment (unlabelled) – assignment (labeled) – free exploration. We again measured **time**, perceived **difficulty**, and **confidence**, analogously to the first study. Again, we let our participants freely

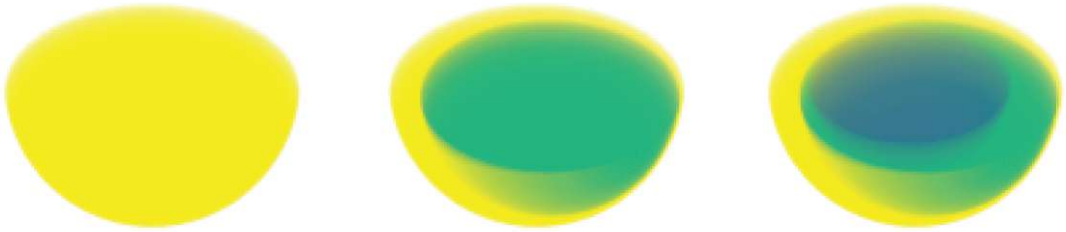

(a) (b) (c)

Figure 5.11: Datasets used for the follow-up study, upper half removed to illustrate composition. Different colors represent different scalar values. (a) Single homogeneous region (sorting and counting task), (b) Two nested regions (counting task), (c) Three nested regions (counting and assignment task).

interact with our study props. The data for each task was collected using a questionnaire. For the **sorting** and **counting** tasks, participants were able to compare three spheres at once and had to assign each sphere to exactly one of three possible answers. The props in the **counting** task were assigned pseudo-randomly for each participant, alternating different combinations of stiffnesses in the two-region and three-region props. For the **assignment** tasks, the participants were handed one sphere at a time according to a balanced, pseudo-randomized scheme. The first three spheres were labeled using a red dot on an indicated position, and then the task was repeated with three different spheres without labels. Finally, we handed the participants the **foot**, **head**, and **fibers** prototypes and briefly explained how they were created. For each prototype, the participants shared their thoughts and findings during the interaction. Lastly, we asked people who participated in the previous study about how their experiences differed in comparison. After each task, we asked the participants for qualitative feedback.

Study Results. For the **sorting**, **counting**, and **assignment** tasks, we report the percentage of totally correct, partially correct, and completely wrong sorting results. For the **assignment** task, we also list these, but calculated per subtask (labeled, unlabeled). We also report time taken, perceived difficulty, and confidence per task. The detailed results are shown in Table 5.3. For the **exploration** task, we code individual insights and list the number of occurrences. We also provide a summary of the received qualitative feedback. The **quantitative** results confirm that all participants managed to **sort** homogeneous squishicalizations correctly, in order of stiffness. In the **counting** task, most participants (56%) were able to identify the number of regions with varying stiffness within the spheres. Five participants confused the sphere with two regions with the sphere with three regions, while two confused the homogeneous sphere with the sphere with two regions. The influence of the labeling of the orientation in the **assignment** task was noticeable, with 78% totally correct assignments in the labeled condition and 43% in the unlabeled condition. In the labeled condition, the innermost region of the props was identified correctly in 77.78% of cases, while the outer and middle regions were identified correctly more often (87.03% and 90.47% times, respectively). For the unlabeled condition, we observe that participants identified the innermost region correctly in 74.07% of the cases, while the outer and middle regions were

Table 5.3: Quantitative study results for the follow-up study. Times in $m:ss$. Accuracy is reported as the number of results with correct sortings (All), the number of results with partially correct sortings (Part), and the number of completely wrong sorting (None). For the Sorting (S) and Counting (C) tasks, the total number of trials is 18 (one per participant); for the assignment tasks, labeled (A*) and unlabeled (A), the total number of trials is 54 (three per participant).

Task	**Accuracy**			**Time**		**Difficulty**		**Confidence**	
	All	Part	None	avg	std	avg	std	avg	std
S	18	0	0	0:14	0:06	1.28	0.57	9.39	1.54
C	10	7	1	01:19	1:05	5.89	1.88	5.61	2.20
A*	42	12	0	02:24	1:03	5.72	2.16	5.67	1.91
A	23	21	10	02:56	1:12	6.78	2.26	4.72	2.22

only identified in 46.3% of the cases.

Our collected **qualitative** data sheds further light on different aspects of the follow-up study. In the free **exploration** task, we counted individual insights per prop. For the foot prototype, 15 participants reported being able to feel the bones inside the structure. 12 participants remarked positively on the realistic feel in terms of stiffness compared to a real foot. We also collected 7 comments on two factors that were perceived as unrealistic: the missing nails and the unnatural surface texture. While the head prototype was supposed to illustrate the position and structure of the brain, only two participants reported feeling a difference between white and grey matter. However, 10 participants noticed a large cavity behind the ear, which is not seen from the outside, and 8 participants found it interesting to explore cavities, such as the nasal cavity and the larynx. In the fibers dataset, 13 participants remarked upon the increased stiffness along the direction of the fibers. Additionally, 4 participants found a difference in elasticity along two axes in the direction parallel to the fibers. 5 participants commented on the visible aspects of the composition of the material's microstructure, such as fiber alignment.

The **sorting** task was considered "easy" or "trivial" by seven participants, but three comments indicate that it is harder to tell apart the medium sphere from the hard than the soft one. For the **counting** task, 8 participants indicated difficulty in identifying a difference between two and three regions. The feedback on the **assignment (labeled)** task shows that four participants found it difficult to identify the inner region and three had trouble identifying region borders. In contrast, for the **assignment (unlabelled)**, five participants found the outer layers harder to determine—with five comments explicitly referring to the missing labeling. Interestingly, three participants found the inner region easier to determine. In **comparison to the first study**, 8 of the 10 returning participants confirmed that the z-axis anisotropy was less noticeable with the new props. In all of the qualitative data, only one statement refers to this effect being noticeable.

5.4 Expert Interviews

To determine potential future usage scenarios for *Squishicalization*, we reached out to domain experts in the fields of extended reality (XR), materials science, and medical education. Every session took 30–60 minutes. The focus was put specifically on a potential application of our proposed workflow as opposed to the general use of data physicalization. We first presented the concept to the experts and then held semi-structured interviews targeting the following questions: *What kind of **data** do you think squishicalizations are useful for? For what kind of **users** are Squishicalizations suited? Which **tasks** could be performed using squishicalizations? Where do you see **weaknesses** in squishicalizations in their current form?* In the interview summaries below, we refer to the domain experts by pseudonyms based on their profession and using their actual pronouns.

Extended Reality. We interviewed on-site a senior XR researcher from our faculty, Dr. X. We brought a prototype version of the foot squishicalization (Figure 5.5a) to explain the concept and let Dr. X freely interact with it. The discussion focused on the simulation of an immersive and realistic environment. As such, the data used in this hypothetical usage scenario would be largely

synthetic. The potential users indicated by Dr. X are professionals in training, as well as XR researchers and consumers. Dr. X identified potential applications for our squishicalizations in **mobile haptics** for VR applications. Here, users enter a VR environment while tactile feedback for specific objects in the virtual space is provided using real-life props. The props are handled by a robot that simulates the behavior of the virtual object. Dr. X added that even though squishicalizations cannot replicate the weight of represented objects accurately, a robot could simulate this by adding force. Furthermore, he indicated possible use in **surgical simulations**. Since squishicalizations can be replicated relatively easily and cheaply using 3D printing, they could be used as destructible props. To provide accurate feedback for such applications, it would be necessary to replicate tactile tissue properties accurately, while a head-mounted display would provide realistic visuals. Lastly, Dr. X proposed that, augmented with pressure sensors, our squishicalizations could serve as **provenance** tools to analyze user interactions. In combination with VR applications, they could provide insights into how people interact with objects in a virtual space.

Materials Science. We subsequently reached out to a visualization expert in the field of materials science, M. This interview was held virtually as the expert is situated in a different city, but the props had been mailed to him in advance. M described the data used in the hypothetical usage scenario as industrial CT scans ranging up to $2000 \times 2000 \times 2000$ voxels. Our method would be suitable for foam-like structures or fiber-reinforced composites. The user group of the applications discussed with M includes material scientists and researchers. As a first possible application, M mentioned **immersive analytics**. Current approaches in materials science visualization use VR in an attempt to enhance 3D spatial perception. M mentioned that physicalization could represent material properties with tactile feedback. Other application scenarios could include the **analysis of fabrication characteristics** of different materials, such as pores and voids, or **defect analysis** and the **temporal analysis of material properties**.

After our initial interview, M provided us with a CT scan from a fiber-reinforced material (the squishicalization of which is shown in Figure 5.5c) and recommended a **follow-up interview** with a fiber-reinforced composite material expert, F. Subsequently, we generated, printed, and shipped the specimen to the two experts for a second virtual interview with both. F commented on the accurate representation of the isotropy in the fiber material. The way the sample cannot be compressed along the fiber direction resembles how such materials may behave in reality, despite the individual fibers being noticeably more elastic than in reality. F found that the compression behavior of the sample in the horizontal plane showed some anisotropy, which could provide further insights into the fiber microstructure. F also mentioned other kinds of **compound materials** that could be interesting to analyze using our approach, such as composite laminates and multi-material components. Both experts indicated limitations for applications in their domain. According to M, some datasets have very small anomalies, which can be undetected or lost during our sampling process. F also mentioned that a more accurate parametrization of the process, for example, based on additional local mechanical properties caused by the microstructure, would help create a more realistic behavior, while larger samples could represent the characteristics more faithfully. To increase realism, they also discussed the possibility of multi-material printing.

Medical Education. We interviewed also Dr. W, a University Professor and specialist in anatomy and anatomical education. He extended the invitation to two medical doctors in training, Dr. K and Dr. A. The interview was held in person in their office and we brought prototypes of the foot dataset (Figure 5.5a) for them to examine. Dr. W mentioned tangible representation of **3D histology** as the first possible application. He described the idea of creating enlarged squishicalizations of cellular structures. Using microscopy data, our method could create a "*tangible cell*", which would be focused on K12 students rather than doctors in training. Additionally, Dr. W inquired about the ability of our method to **display large cavities** in human bodies, in a way that their proportion and position are understandable. He also proposed creating waterproof models of these cavities to **illustrate blood flow** through cardiac volume. As doctors in training, Dr. K and Dr. A identified a potential for creating **tangible representations of rare conditions** for anatomical education. Doctors use palpation to detect anomalies in tissue using their sense of touch. Some of these conditions do not appear regularly in cadavers and are exemplified through prepared specimens. Dr. A explained that the tactile feel of prepared specimens differs from live examples, while Dr. K added that such tactile stimuli can not be taught in books. Dr. W pointed out the importance of accurate tissue reproduction for the method to be usable for professionals in training. He stated an accurate mapping from tissue type to elasticity would be needed. Additionally, Dr. K pointed out that the surface properties of the presented squishicalizations were unrealistic.

5.5 Takeaways and Future Directions

Data physicalization research leverages the tactile properties of physical sculptures. As opposed to other recent works that use surface properties of 3D printed objects to express characteristics in data [DMSJ23, PRC+23, MSF24] or other fabrication solutions with [BFY+24, TRBA20] or without the possibility to measure user input [SSJ+14], with *Squishicalization*, **we express the characteristics of data in physicalizations beyond the surface**. We encode scalar values as local densities in printable artifacts to add a tactile dimension to the generated data physicalizations. Moreover, our *Squishicalization* pipeline is **designed for use with any consumer-grade 3D printing that can make use of elastic materials**. All tools and materials are available and affordable for home use, and no manual labor is needed for assembly.

Performance. Compared to other microstructure-based approaches [SBR+15,MHSL18], *Squishicalization* is **applicable to volumetric data without prior segmentation**. Similar to volume rendering, we employ piecewise linear DTFs to map local density to scalar intervals. In terms of performance, Martinez et al. [MHSL18] name extraction of region borders as the biggest bottleneck in their Voronoi tessellation-based approach. In *Squishicalization*, we eliminate the performance impact of the tesselation by using labeled volume isosurface extraction [Fri22]. Parallelization of the region extraction part on the GPU and restricting tessellation to regions of interest increases our performance, as we showed in Section 5.3.1. Compared to Lu et al. [LSZ+14], who demonstrate their approach with just 100 points, our method scales well with a much higher number of seed points (see Table 5.1).

Parameter Space Exploration. *Squishicalization* has a vast parameter space to adjust for material and fabrication-related variations. In Section 5.3.3, we document the results of our mechanical tests on samples with different densities and variations in wall thickness. The axial load tests show that **our method can reproduce a wide range of elasticity**. We also show that the samples have near isotropic behavior in the horizontal plane. As highlighted later by some experts in our interviews (Section 5.4), a more accurate mapping between scalar values and elasticity is needed to accurately reproduce desired behavior. Currently, we only use voxel distance mapping scalar values to construct our DTFs. **A first step to more accurate representations could be to incorporate voxel spacing and point densities into DTFs**. Determining the precise relationship between sample density and elasticity for a given wall thickness and material is essential for this. Our mechanical and perceptual evaluations point towards a nonlinear relationship between these factors. Combined with traditional dart-throwing Poisson-disk sampling approaches [Bri07], precision can increase at the expense of control over the currently user-selected amount of samples.

Suitability for Selected Applications. Our results show that our prototypes could stimulate the tactile sense of our participants to determine the size and positions of medium and large-sized heterogeneous regions, as well as nested components. We also show that the elastic properties of human tissue can be emulated realistically through our mechanical evaluation. This indicates that our technique could be, for instance, useful as a **training tool for palpitation** of rare conditions for medical professionals and demonstration of self-palpitation of irregularities for laypeople. Also, our interviews with material scientists revealed that our method could be useful to **analyze and illustrate material characteristics** of compound materials. Extending the method to provide tactile input measurement, similar to Bae et al. [BFY+24] and Tejada et al. [TRBA20], could further support the analysis of users' **tactile interactions**. Tangible encoding of **uncertainty** in scalar data could be encoded in stiffness, e.g., with harder structures indicating higher certainty, and softer structures the opposite.

Oppositely, as study participants struggled to determine the directions of gradients, we have indications that *squishicalizations* might **not be suitable** for identifying the structure of more complex mathematical formulations, such as probability distributions. Our technique produces rather rough surfaces, which could confuse sightless people in **accessibility scenarios**, which could be improved in the future with an additional sensory channel. Finally, several participants stated that using tactile perception is **unusual**. During our study, participants were sometimes concerned they could destroy the props, or felt inhibited in applying too much force. This indicates that "hidden" direct interactivity, such as exploring the sub-surface properties produced by *Squishicalization*, needs to be encouraged to be fully leveraged.

Refinements. The initial perceptual study and the mechanical evaluation used samples with a wall thickness of $0.3mm$. This way we could produce a large number of samples in reduced time and with moderate material use. However, the axial load tests indicate that this wall strength increases the stiffness of the samples in the printing direction. The low vertical resolution of these prints leads to disconnected staircase artifacts between the layers, resulting from the slicing method. Combined with the additional walls, bottom, and top cover of the cubes used in our

user study, we see a large impact on tactile perception, reflected by the qualitative comments of our participants and also their low performance. In our follow-up study, we reduced the wall thickness of the props and used a spherical shape without stiff outer surfaces, guided by the mechanical evaluation results. **Using wall thicknesses around** $0.2mm$ **yielded more accurate results and led to less noticeable anisotropy.** This only requires an increase of sample points to reach comparable levels of stiffness. Staircase artifacts could be reduced by increasing the vertical resolution of the prints, or by refining the slicing to allow for horizontal walls. Also, conical slicing [WGEJ21], as opposed to vertical slicing, could be used to create more complex squishicalizations, even with overhanging structures.

Conclusion. With *Squishicalization* we show a method to transform data into tangible representations using consumer-grade 3D printing technology. Using commercially available elastic 3D printer filaments, we create and explore tangible representations of 3D volumes that encode scalars into local deformation behavior. Our results show that users are able to perceive nested objects composed of layers with varying elasticity successfully, and expert interviews indicate use cases for our technique in many different fields. *Squishicalization* represents a first step towards tactile representations that leverage human sensing capabilities in an innovative way.

CHAPTER 6

Heart Machine: Exploring an Intermodal Physicalization of a Dynamic Physiological Process

This chapter is based on the following publication:

> Pahr, D., Ehlers, H., Wu, H.-Y., Waldner, M.,Raidou, R.G. (2024). **Investigating the Effect of Operation Mode and Manifestation on Physicalizations of Dynamic Processes**. In *Computer Graphics Forum 43 (3)*. DOI: https://doi.org/10.1111/cgf.15106

Related work on the integration of the observer as an active part of physical data representations focuses on the illustration of static datasets or facts, whereas dynamic processes remain largely unexplored. We propose a study as a first step to consider the manifestation of a representation and the way an observer interacts with it as independent factors. In doing so, we avoid the confounding of beneficial effects of either factor. Furthermore, we individually assess how these factors influence not only learning but also task load and engagement. Our manually operated representation encodes information by integrating the actions of a user into the representation, by intermodally communicating via visual stimuli and tactile feedback from physical interaction (Figure 1.2d).

6.1 Introduction

Visual representations, whether in virtual or physical form, are integral tools in education, widely embraced for their effectiveness. In STEM fields, the use of visual representations has been extensively studied with concepts like pulleys in physics [GCC+10, FH95, BS09], prompting an exploration into the factors contributing to their widespread adoption. Among others, interaction

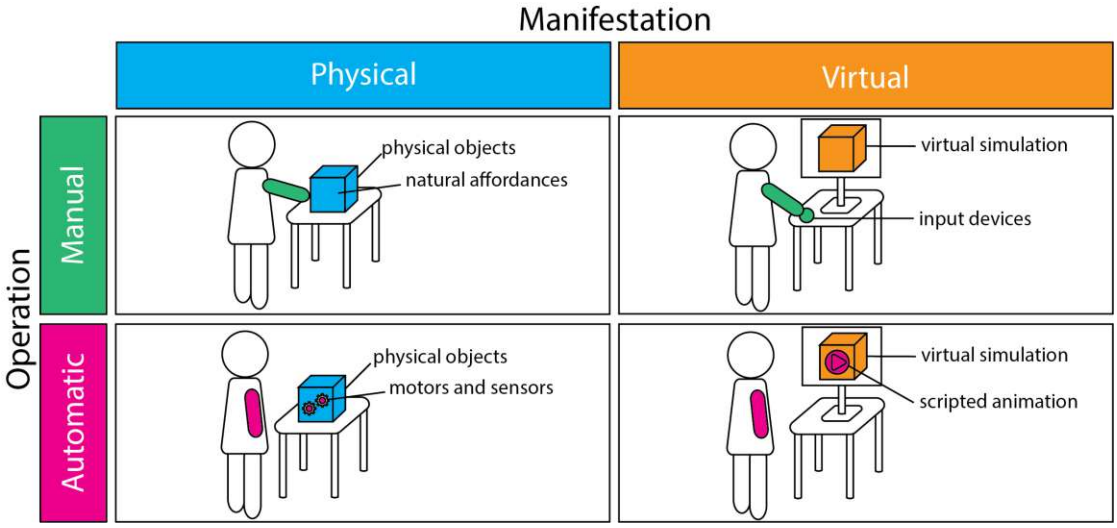

Figure 6.1: The two-dimensional design space explored in our study delves into the communication of complex dynamic processes. The two dimensions comprise the *manifestation* of a representation (physical or virtual) and the mode of *operation* (manual or automatic).

with these representations is one of the most documented elements as it enhances the audience's understanding.

In the medical context, conveying (patho)physiological processes poses unique challenges compared to more straightforward concepts in physics, such as pulleys. For instance, communicating the process behind a "healthy" cardiac cycle to laypeople involves explaining the coordinated phases of systole and diastole, the role of heart valves and chambers, and the efficient circulation of blood [WG10]. Communicating potential pathophysiologies, i.e., functional changes accompanying a pathological condition, adds complexity as it involves comprehending the impact of disruptions on heart function and overall health. Yet, patient communication and education are crucial for informed decision-making, active participation in healthcare, and an enhanced overall patient understanding of their conditions and treatment options [MGS+21].

Within the realm of visual representations, physicalizations emerge as a unique opportunity for bringing data into the physical space [DJVM20]. This is primarily due to the distinct characteristics of physical representations, such as the use of physical embodiment and natural affordances to convey meaning and engage an audience [ZVM08]. While physicalizations find widespread use in STEM [GCC+10], their application in explaining dynamic processes of higher complexity remains relatively unexplored [Rau20].

Recent work in data physicalization targeted the creation of engaging and low-cost anatomical models for edutainment and patient education—without addressing pathophysiological or any other dynamical aspects [SB17, PWR21, SKRW22]. All prior examples employ *indirect interaction* that acts similar to interface controls [BZW+22]. Conversely, multisensory displays [HH16] use *direct interaction*, directly stimulating a user's sensory affordances [BZW+22]. The concept of direct interaction opens up additional channels—for example, the use of kinaesthetics—to

encode data into physical activity performed by observers [HMC+20]. Directly interactive models hold promise as representations of pathophysiological processes, but they remain largely unexplored and unassessed in this context.

Research in data physicalization has sought to quantify the value of physical data representations [JDF13, SSB15]. This has often been done using comparative methods, where physical data representations are compared against similar screen-based ones. Until recently, evaluations of both physical and virtual data representations focused primarily on efficiency, measuring how quickly insights can be derived. This narrow focus has inadvertently left unexplored—or even obscured—potential advantages of data physicalization, especially for the medical domain where audience engagement in the communication of complex processes is crucial [MGS+21]. Recently, new metrics with a focus different from efficiency or comprehension have emerged offering to support the evaluation of such concepts [WSK+19].

To fill this gap in research, we investigate the potential of directly interactive models that represent dynamic pathophysiological processes. We also delve into assessing the implications and benefits of data physicalization beyond traditional efficiency metrics. To this end, we explore the effects of **manifestation** and mode of **operation** of a representation on the communication of pathophysiological processes on **observer understanding (Q1)**, **subjective task load (Q2)**, and **enjoyment (Q3)**. We consider the manifestation—virtual or physical—and the mode of operation—manual or automatic—as two individual factors in a full factorial study with 28 participants. Figure 6.1 shows an overview of the resulting design space. We employ a **mixed methods** study design. First, we use quantitative methods to validate a-priori-postulated hypotheses about understanding, task load, and emotional engagement. Subsequently, we analyze the qualitative feedback of our study subjects. Finally, using triangulation [OMN10] we gather the insights obtained through both methods, to draw further conclusions on our findings.

The contribution of our work stems from the results of an exploration of a two-dimensional design space for educational models for pathophysiological processes (Figure 6.1). We confirm that both physical manifestation and manual operation increase engagement. Qualitative feedback reveals that manual interaction with physical representations augments non-visual stimuli, leading to subjective knowledge increase. With the results of this study, we discuss the implications that arise from our design space in educational pathophysiological process representations.

6.2 Methodological Approach

Design Space. In this work, we investigate the effects of the **operation** mode and **manifestation** as separate factors within physicalizations of dynamic processes, such as those present in pathophysiology. To this end, our design space is built upon two dimensions: the *mode of operation*, describing the way the representation is operated, and *manifestation*, describing the medium in which the representation exists. An observer operates the representation in a manual way or the representation is automated. The manifestation of the representation can be either physical or virtual. Our two-dimensional design space for process physicalizations is schematically represented in Figure 6.1.

Representation-based learning, as well as data physicalization, has previously investigated **physical manifestation** as a central concept. In the physical manual quadrant of our design space we consider only data physicalizations with **direct interactivity** [BZW+22]. Examples, among the related works, are INTUIT [DMSJ23] or "Move&Find" [HMC+20]. In the virtual manual quadrant, direct interaction is relegated to the input and output devices of, e.g., a PC. We choose not to add additional indirect interaction modalities to our virtual representations. Examples of indirect interactions from the related works are Boucheix and Schneider's [BS09] controlled animations, or the visualizations used by Jansen et al. [JDF13] in comparison to their physical bar chart. In the automated quadrants, we remove the interaction with the representations completely. An example of a virtual automatic representation is Hurtienne et al.'s [HMC+20] animated cyclist, or other animated representations [BS09]. Oppositely, Perovich et al.'s [PWB21] floating lanterns representing environmental pollution are a physical automatic physicalization.

Research Questions and Hypotheses. While prior work has not shown that physical representations are more efficient for information retrieval [JDF13, SSB15, HMC+20], this has not been yet evaluated for representations of dynamic processes. Increased physical load interacting with physical representations has been observed in the past [HMC+20]. Finally, physical representations have been shown to be more engaging [HMC+20] than virtual ones. We, therefore, pose three separate research questions regarding **understanding (Q1)**, **task load (Q2)**, and **enjoyment (Q3)** of pathophysiological process representations—specifically, w.r.t. how they are impacted by different modes of operation and manifestation:

H1: Physical and manually operated representations are more effective in conveying processes than virtual and automated ones.

H2: Physical and manually operated representations entail a higher subjective task load than virtual and automated ones.

H3: Physical and manually operated representations are more enjoyable than virtual and automated ones.

Metaphors [ZVM08] and visual abstractions [VCI20] in the context of data communication serve to simplify intricate physiological processes, helping both professionals and the general public grasp complex concepts. To enhance comprehension for a lay audience, we introduce an abstract metaphor that elucidates a complex process within the human body: cardiac function. We discuss our metaphor choice and the respective representation design arising from it in Sections 3.1–2. Deriving from the design space, we create four representations, one for each quadrant. These are physical manual (**PM**), physical automated (**PA**), virtual manual (**VM**), and virtual automated (**VA**).

Subsequently, to investigate our previously formulated hypotheses with regard to the conceptualized metaphor and its corresponding representation design, we employ both quantitative and qualitative methods. We initially look at how the mode of operation and manifestation influence an observer's understanding of a process representation (**H1**). We also measure individual task load (**H2**) and enjoyment, in terms of emotional engagement (**H3**), in interacting with said

representations. Subjective preference will be used as an additional measure for enjoyment (**H3**). Written statements collected during the study will serve as the basis for our qualitative evaluation.

6.2.1 Modeling Basic Cardiac Function with a Metaphor

A primary goal for (physical) data representation is the communication of insights to laypeople with limited domain knowledge. In the context of data visualization and data physicalization alike, both metaphors and visual abstractions serve as essential tools for conveying complex information to diverse audiences. Metaphors in data physicalization function as powerful cognitive tools, enabling observers to relate to the representation of intricate processes and complex data sets by leveraging familiar concepts [ZVM08]. Visual abstractions, on the other side, filter out unnecessary details in data representations allowing users to focus on crucial elements [VCI20].

We introduce an abstract metaphor to elucidate the complexities of the physiological processes of cardiac function. Cardiovascular diseases are the leading cause of global mortality, while many of the underlying risks can be addressed through behavioral changes [Rot20]. Campaigns [SFJ+22] and installations [McA16] communicating cardiovascular function and associated risks are widely employed to raise awareness in the general population. Still, data physicalization approaches looking in this direction are scarce [ASS+19] but are anticipated to facilitate a tangible and accessible comprehension of cardiovascular processes by transforming complex data into interactive and visually intuitive representations. To ensure broad applicability, our model is designed for both virtual and physical representation, accommodating automated and manual operation. This versatility enhances the accessibility of medical information for laypeople, aligning with our goal of simplifying complex concepts in health education.

The Pump Metaphor. In the circulatory system, the function of the heart is to pump blood through the circulatory system. During the diastole, the cardiac muscle relaxes and the chambers fill with blood. In the systolic phase, the muscle contracts, expelling the blood to the lungs and the peripheral vessels. Besides the actual meaning that the heart transports fluids, this also provides us with an intuitive metaphor for its function. Pumps are a familiar concept, inherently understandable to people. Beyond the aforementioned metaphorical value of the subject, Offenhuber's perspective [Off20] gives additional directions. The pump metaphor is very close to the actual function of the heart, defining the function in an *ontological* sense. Laypeople may also not be familiar with the specific variables that medical professionals use when they speak about bodily functions, like blood oxygenation, blood pressure, or heart-stroke volume. We argue that a *relational* perspective is better suited for non-experts. Deconstructing the heart function with the pump metaphor allows for a deeper understanding of the regulation process. Using an ontological-relational perspective limits the use of numeric representations with which a lay audience might not be familiar.

Representation Design. We now proceed to design our models around the pump metaphor for the human heart. As mentioned above, the basic heart function is the transport of blood through the circulatory system. Physically, we choose to represent the human heart with a commercially available pump, suited to transport the same amount of liquid. The physical pump is operated

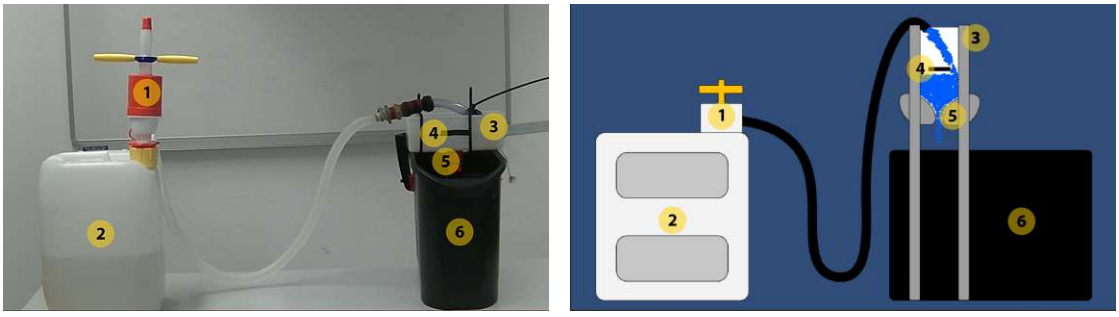

(a) (b)

Figure 6.2: Comparison of the physical (a) and virtual (b) representation. The pump (1) on the canister (2) represents the basic heart function. Liquid is transported to the buffer vessel (3) on the right. The level indicator (4) marks the point at which the water level has to be kept to represent 5–6 liters per minute of flow. The valve (5) leads into the reservoir (6).

by moving a piston up and down, which first creates negative pressure in the inlet when it is moved upwards. This is similar to the *diastole* of the cardiac cycle. When the piston is moved downwards, the pressure inside the pump, towards the outlet is increased. This resembles the *systole*. The demand for blood in the body is represented by a buffer that is drained at a constant rate. This rate resembles the pump output of the heart. To keep the buffer filled, constant operation of the pump is required. In this state, the system represents the normal function of a heart. This manually operated, physical representation (**PM**) is shown in Figure 6.2a.

To compare manual and automated representations, we ensure that operating the pump can be taken over by a set of sensors and actors. While it would be preferable to leave the manual pump in place and design a system for automatic operation, we opted for a simpler design in favor of practicality. We utilize an electric pump operating similarly to the manual one. Electrical pumps are typically designed for continuous operation as opposed to our chosen manual exemplary. However, this would destroy the systole/diastole metaphor. Therefore, we split the pump into one simulated and one operational part. The simulated part consists of a motor, operating a piston complete with simulated handles in reciprocating motion, to imitate the operation of a user. The operational part consists of the electrical pump, switched on and off in the rhythm of the simulated pumping. The system is controlled by a level sensor in the buffer vessel, which switches on the pump when the water level in the buffer vessel is under a certain threshold. Thus, we obtain a physical representation that works autonomously (**PA**).

After designing the physical representation, we now reproduce the pump metaphor in a virtual setting. To keep the interaction with the model similar to the physical installation, we simulate the process using Unity [Tec], since it provides a basic physics simulation environment that serves our purposes. For the simulation, we take multiple considerations into account. It could be done in a 3D environment, where interaction is still performed via mouse and keyboard while introducing additional complexity for observers. Virtual and augmented reality simulations of the process could be used as well. These technologies present additional complications that could influence our metrics. We opt for a 2D representation to limit interaction with the virtual representation to

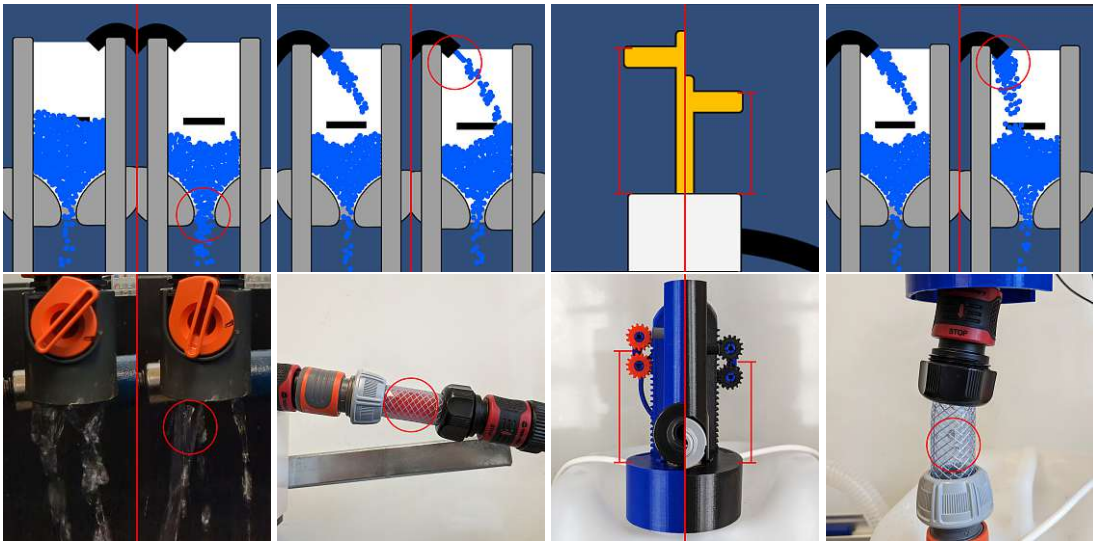

(a) (b) (c) (d)

Figure 6.3: Scenarios, representing different heart conditions with a virtual (top) or physical metaphor (bottom). (a) Exertion: Higher oxygen demand is represented by opening the drainage valve of the buffer vessel. (b) Aortic stenosis. The narrowing of the aortic valve is represented by a narrowing of the pump's outlet. (c) Systolic heart failure with reduced stroke volume: Limited heart volume is represented by limiting the movement range of the piston. (d) Cor Pulmonale. Difficulty in transporting blood to the lungs is simulated by puncturing the inlet of the pump.

the same interaction points that exist in the physical version, simulated with keyboard and mouse interactions to provide a typical PC environment. In both representations, the handle of the pump is moved up and down to collect and expel liquid, and the targeted flow rate is symbolized by the watermark in the draining vessel. This way we limit the focus of our study to the factors of manifestation and mode of operation.

In our Unity simulation, we show a schematic view of the physical installation. The canister, pump, and buffer are represented by 2D objects, as shown in Figure 6.2b. The piston is in the same position, on top of the canister, operated by moving the piston up and down. An observer can move the piston with mouse and keyboard gestures, while the drag and weight of the object are simulated in the physics engine. When the piston is moved downwards, the ejected water is represented by particles exiting the hose into the buffer. The particles exit the buffer from the bottom, depending on how wide an adjusted opening is, simulating the draining of the buffer. This represents the process virtually, with a manual operation (**VM**). These are the same interaction points as in the physical installation. The motion of the piston can be automated using a simple script to obtain an automated virtual representation (**VA**), as the virtual counterpart of the automated pump.

6.2.2 Variations of the Model

For the representations of the cardiac function, as shown in Figure 6.2, we now introduce different physiological and pathological states to create variations of the model. An overview of these scenarios and how they are represented in the physical and virtual representation respectively is shown in Figure 6.3. The individual scenarios illustrate how cardiac function changes in contrast to normal function. We use examples of common pathologies that can be illustrated using our model. In a consultation with a medical doctor, we ensured that the scenarios were depicted correctly, within the limited scope of our metaphor. In our study, participants are presented with each of the four variations of our representation (**PM**, **PA**, **VM**, **VA**), each depicting a different scenario. Their task is to determine the difference between normal cardiac function and the altered state in the specific scenario. The four scenarios (**S1–4**), are discussed below:

Exertion (S1) leads to increased demand for oxygen in the body. While at rest the cardiac output is $5-6\ L/min$, this output can increase drastically during exercise, to even more than $35\ L/min$ in elite athletes [KL22]. In our model, we represent the cardiac output by a slowly draining buffer vessel. In the interactive version, observers have to keep the buffer vessel filled around a marked threshold. In the automated version, a binary level sensor determines the fill state and pauses the pump when the sensor triggers. A state of exertion is simulated by opening the drainage valve of the buffer vessel further, as compared to the normal state. This is illustrated in Figure 6.3a.

Aortic Stenosis (S2) is caused by a stiffening of the aortic valve, leading it to not properly open. The heart has to work harder to contract the left ventricle and the systolic pressure increases [PA22]. To represent this physically, we 3D printed an obstruction of the pump's outlet, reducing the area of the outlet hose by half. In the virtual representation, this is simulated by increasing the resistance of the piston, as well as increasing the speed of the particles spilling into the reservoir. We illustrate the visible differences of this in Figure 6.3b.

Systolic Heart Failure (S3) or heart failure with reduced ejection fraction is a condition where the left ventricle's pumping output is reduced. Here, the heart's ability to contract is limited by physical factors [HL22]. We represent the reduced pumping capability by limiting the movement range of the piston in the physical model. We automate the physical pumping process by simulating a reciprocating pump. The reduction of stroke volume is achieved by reducing the length of the rotating lever operated by the motor. Virtually this is represented by limiting the range of the piston. This can be seen in Figure 6.3c.

Cor Pulmonale (S4) is the fourth condition we simulate. In this case, the part of the heart that pumps blood into the pulmonary artery is affected. Depending on the cause, this does not entail a reduced ejection fraction [VQL+06]. We represent pulmonary hypertension by drilling a hole into the inlet of the pumping mechanism. In the virtual representation, we simulate this by a reduction of the number of particles spawned by the outlet and a decreased weight of the piston. Figure 6.3d depicts this condition.

6.3 Study

We use a 2×2 full-factorial within-subject design, with manifestation and mode of operation as independent variables. Manifestation can be physical (**P**) or virtual (**V**), and operation can be manual (**M**) or automatic (**A**). This results in four representations with different combinations of manifestation and mode of operation (**PM**, **PA**, **VM**, **VA**), which are illustrated in Figure 6.1. The scenarios were assigned to representations using a 4×4 Graeco-Latin Square design [Moo16] to combat familiarization effects. In a sequence of four runs of the study, the order of scenarios, as well as the order of representations was unique. A number of participants (28) that is a multiple of four ensures that each combination of model and scenario is examined an equal amount of times.

6.3.1 Participants

We recruited 28 participants for our study from staff and students in the Faculty of Informatics at TU Wien, Austria. 20 identified as male, and eight as female. 16 of them used visual aids. Age was collected in discrete intervals between 18 and 54 years. The largest age group was between 25 and 34 years old, which was 20 participants in total. Two participants self-reported to be well-versed in anatomy, 14 participants reported basic knowledge, 11 reported themselves to be novices, and two participants reported no anatomical prior knowledge. One participant has a high school degree, 11 have a bachelor's degree, 12 have a master's degree, and four have completed a Ph.D. level education. All trials of our study took place in person, in the same environment. The participation was voluntary, and no compensation was provided. Participants could opt out at any point if they felt discomfort.

6.3.2 Procedure

Every participant interacted with one representation (**PM**, **PA**, **VM**, **VA**) per scenario (**S1–S4**), within a randomized order per individual subject. Each run of the trial started with the participants being led into the room where the experiments were conducted. The participants were handed a written introduction summarizing the nature of the tasks they were about to be presented with. We explained the procedure and clarified essential vocabulary. Subsequently, the participants were asked to read an introduction to the first scenario and we explained the task.

For each scenario, participants were first presented with a representation of normal heart operation, to familiarize themselves with the model and how it represents normal heart function. They were allowed as much time as they desired with the normal functioning model. After that, we introduced the complication into the model, i.e., the variations discussed in Section 6.2.2. The participants were allowed as much time as they wanted with the altered state. Then they were asked to complete a quiz, to determine how much they understood a complication (**H1**). After the quiz, we asked them to complete two additional questionnaires.

Task. The task was to answer questions about how the different complications affect certain vital parameters in a multiple-choice quiz. The participants were presented with one representation, first in a normal state, then in a complicated state. Then they had to decide how the vital

parameters **heart rate**, **cardiac preload**, and **afterload** were affected in the scenarios. The multiple-choice quiz allowed the answers **increased**, **lowered**, and **unaffected**. The quiz also contained binary questions about the presence of the symptoms **fatigue** and **shortness of breath**. We allowed the participants to indicate they are **unsure** about the answers, to prevent them from making random choices. These parameters were the same for every scenario and were introduced and thoroughly explained to subjects beforehand.

Measuring Task Load. The first questionnaire the participants were confronted with was the NASA Task Load Index (TLX) [HS88]. We closely kept to the instructions provided by NASA with only one exception: We used a discrete scale from 0 to 10 for the questionnaire as opposed to the 0–100 scale with 21 increments that NASA suggests. The TLX uses 6 components to measure the subject perceived load during a task. The individual components are mental demand, physical demand, temporal demand, performance, effort, and frustration. These individual factors are measured on a scale from high to low (except for performance which is measured from good to bad). We used the resulting task load index to test for **H2**. During the first experiment of every trial, we performed the scale ranking procedure for the TLX. This procedure is used to identify participants' subjective biases for the different components of the task load index. The task load index for a given task was then calculated by computing the weighted average of the participant's ratings.

Measuring Emotional Impact. The second questionnaire assesses emotional engagement, as proposed by Wang et al. [WSK+19]. They propose to measure engagement as comprised of the following categories: creativity, affective engagement, physical engagement, intellectual engagement, and social engagement. The items of the questionnaire consist of statements the participants rate on a scale of 0 to 10, to be consistent with our implementation of the NASA-TLX. To keep the questionnaire reasonably concise we used a single question for each factor of engagement as proposed by Wang et al. Similar to the TLX, we computed the engagement score as the average of all components, however, we performed no scale-weighing procedure for our engagement score. This score was used to test for **H3**.

Additionally, after every experiment, we asked the participants to leave **written feedback**. While the participant fills the questionnaires we prepared the next experiment. The process was repeated until the participants had seen all four representations and filled out the corresponding questionnaires. Subsequently, the subjects were asked to fill out a final questionnaire about their overall experience. In it, they are asked to rank the models from most to least enjoyed. This ranking was used in addition to the engagement score to test for **H3**. The final questionnaire also contained two **open text questions** about the most and least like experiences with any representation overall. The written statements were used for the qualitative analysis.

Subjects were provided instructions for both parts of the questionnaire, with written explanations of every item, which they could refer to at any point. At the end of the session, there was a debrief, where we asked the participants about their overall thoughts on their experience with the models, and thanked them for their participation.

6.3.3 Analysis

We used a mixed methods approach to address our three hypotheses. This section describes the analysis performed on both quantitative and qualitative data.

Quantitative Analysis. We collected the results of the quiz (**H1**), NASA TLX (**H2**), and our custom engagement questionnaire (**H3**) for all experiments of all study participants. To analyze the influence of these two factors on the variables quiz-score, TLX-score, and embodiment score, we opt to employ Wobbrock et al.'s Aligned Rank Transformed ANOVA (ART-ANOVA) [WFGH11]; a non-parametric framework reliant on an initial rank transformation of the data that allows for the subsequent use of complete two-way ANOVA models. This non-parametric model was used because the model's assumption of normally distributed residuals was rejected when probed with Shapiro-Wilk tests. The statistical significance of the factors was evaluated for each such model using ANOVA-standard t-tests. Multiplicity was accounted for by adjusting the significance threshold using a standard Bonferroni correction. For **H3** we additionally investigate participant's preference in a model's manifestation and mode of operation. Participants ranked their personal preferences for each of the four combinations of operation mode and manifestation. These ranks, owing to their obvious non-normal distribution, were tackled non-parametrically. More specifically, an initial omnibus Friedman test [Fri37] was performed to investigate whether any of the four combinations were statistically significantly different from each other, in which case it was followed by a series of Wilcoxon Signed-Rank Tests [Woo08] to probe their pairwise differences. For these pairwise tests, multiplicity was again accounted for using a Bonferroni correction of the significance threshold across the six comparisons performed.

Qualitative Analysis. We used a coding approach to quantify our collected qualitative data [Chi97], with **three independent coders**. For this, we broke down the collected written feedback in our questionnaires into 78 utterances. In our qualitative analysis, we performed both deductive and inductive coding. As a first step, we assigned each comment to the representations it referred to. This was decided based on the context given by the questionnaires and the content of the statement. Multiple assigned representations were possible, for example, if a statement read "*... in the physical variants...*", we assigned it to both of the physical representations. We found some statements referring exclusively to our used metaphor, which would apply to every representation. These will not be used to validate the hypotheses and will be discussed separately. In the **deductive coding** step, we had three individual coders perform an assignment of the statements to the aspects of understanding (**H1**), task load (**H2**), and subjective enjoyment (**H3**) while rating them as positive or negative. Monitoring the Krippendorf α value [Kri18] for inter-coder agreement, we iteratively refined the assignment until an α greater than 0.8 was reached. The disputed cases were then decided by a majority vote. Finally, each coder performed an **inductive coding**, assigning common concepts to their choices. These codes were unified in a final session. We use this categorization to draw further insights from the feedback.

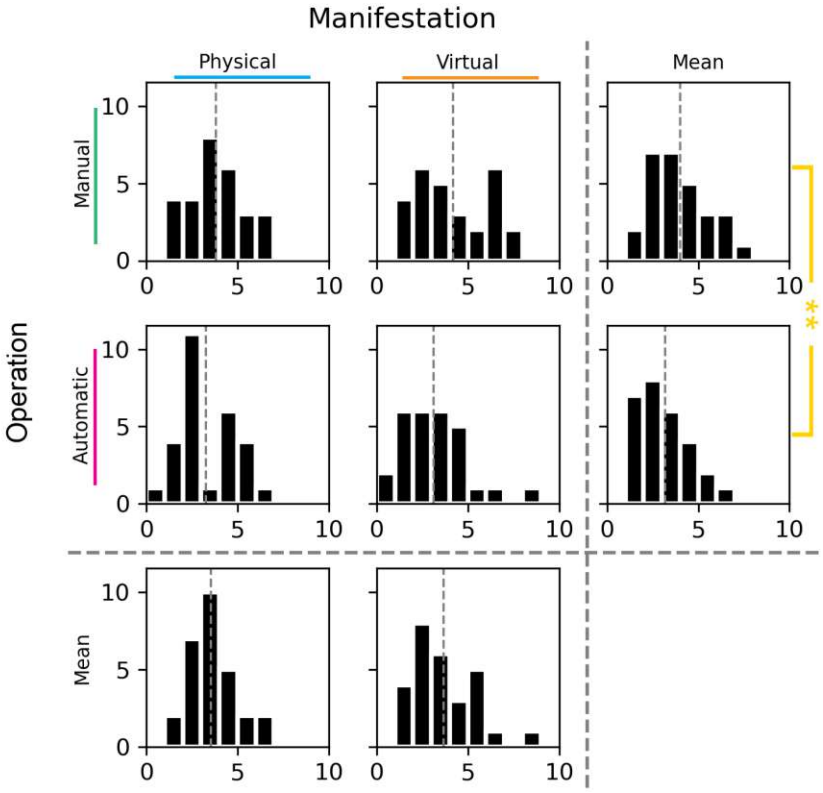

Figure 6.4: Histogram matrix of TLX scores by manifestation and operation mode. Average scores per participant across dimensions are shown behind the dashed lines. The manual row shows an increased task load.

6.4 Results

In this section, we present the findings of both the quantitative and qualitative parts of our study. We then draw our conclusions in relation to each of our hypotheses in a triangulation.

6.4.1 Quantitative Results

Understanding (H1). The quiz scores yielded diverse results across both manifestations and interaction modes. The ART-ANOVA conducted showed *neither a main effect for manifestation* ($F(1,27) = 0.025; p = 0.87$) *nor for operation* ($F(1,27) = 1.28; p = 0.27$). We could also not show an interaction effect between operation and manifestation to be statistically significant ($F(1,27) = 0.01; p = 0.92$).

Task Load (H2). The weighted NASA-TLX scores are shown in Figure 6.4. Firstly, the ANOVA showed no significant effect of manifestation on task load ($F(1,27) < 0.01; p = 0.99$). However, the ANOVA highlighted a *significant effect of manual interaction on task load* ($F(1,27) <$

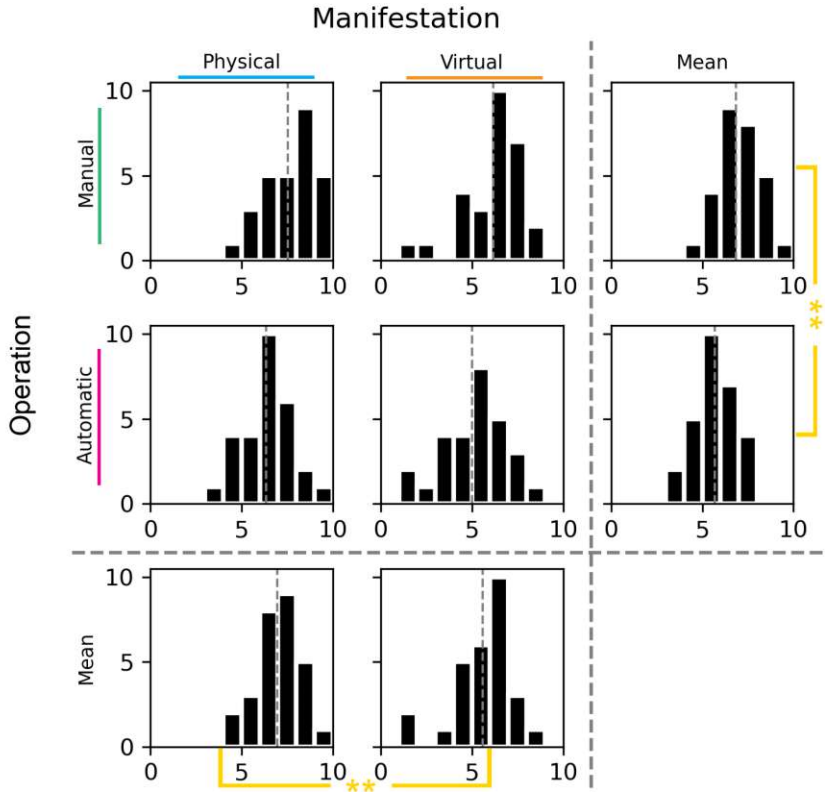

Figure 6.5: Histogram matrix of engagement scores by manifestation and operation mode. Average scores per participant across dimensions are shown behind the dashed lines. Both manual and physical sides show an increase in engagement.

7.94; $p < 0.01$). Participants, as expected, reported higher physical demand in the interaction with physical representations (median=3) compared to virtual ones (median=1.5). When comparing task load for different operation modes, manual models have a notably higher physical demand component (physical demand median=4.5) than automated ones (physical demand median=0). The effort component was also consistently rated higher for manual models (median=5) than for automated ones (median=3.25).

Emotional Engagement and Preference (H3). Figure 6.5 shows the results of the engagement questionnaire. Here, the conducted ANOVA showed a *statistically significant effect of manifestation* on the engagement score ($F(1,27) < 18.73; p < 0.01$). The median engagement score for physical models is 7, while for virtual ones it is 5.95. The *effect of manual interaction on engagement was also statistically significant* ($F(1,27) < 31.61; p < 0.01$), with the overall score for manually operated representations (6.95) higher than that of automatic ones (5.75).

When comparing engagement scores of physical and virtual representations in detail, physical models have produced higher scores across several engagement components. More specifically,

creativity, affective engagement (medians 6.5 vs 5), social engagement (medians 7.5 vs 6.75), as well as physical engagement (medians 6 vs 4.25) all receive higher median ratings for physical representations. Comparing engagement scores between the two modes of interaction, manual representations receive a noticeably higher physical (medians 7.5 vs 4), social (medians 7.75 vs 6), and affective (medians 6.25 vs 5) engagement score. We also did not find a significant correlation between task load and engagement.

We finally show the results of the ranking of the four representations in Figure 6.6. To probe this hypothesis quantitatively, we performed an omnibus Friedman Rank Sum Test to compare the distribution of ranks across the four combinations of manifestation and operation mode ($\chi^2 = 41.571$, $p < 0.01$). Given the Friedman test's significance, we subsequently performed paired, post-hoc Wilcoxon Rank Signed Tests to investigate each of the six pairwise differences between models. We observe significant differences between most variants, with the notable exception of physical automated and virtual manual representations ($p = 1$).

6.4.2 Qualitative Results

We present a summary of the result of the deductive coding step in Figure 6.7. The codes that emerged in the inductive coding are explained below, in conjunction with the associated deductive codes. Selected examples for statements are given in quotes.

Understanding (H1). We found 16 statements that described increased understanding when using physical manual representation, but only two negative statements on this. We registered six statements reporting *non-visual insight*, i.e. insights not gained over the visual channel, five of which directly referenced added tactile information ("*[i liked the] haptic feedback of the pressure on the pump*"), and one referring to audible changes in the system. Additionally, six statements reported *visual insight* gained from physical manual representations ("*seeing the flow rate as water gave me a good idea of how the volume moved*"). Conversely, the automated representations with physical manifestation received five and virtually manifested two

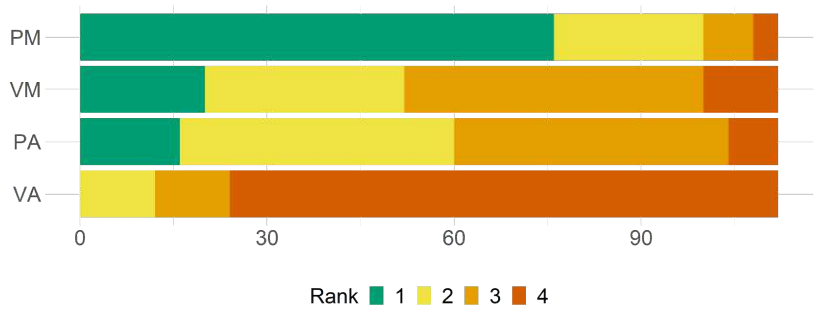

Figure 6.6: Stacked bar charts of preference votes for the different models. All participants chose their most (green) to least preferred (red) representation in the final questionnaire. Physical manual representations are significantly more preferred than others, while virtual automated are the least popular.

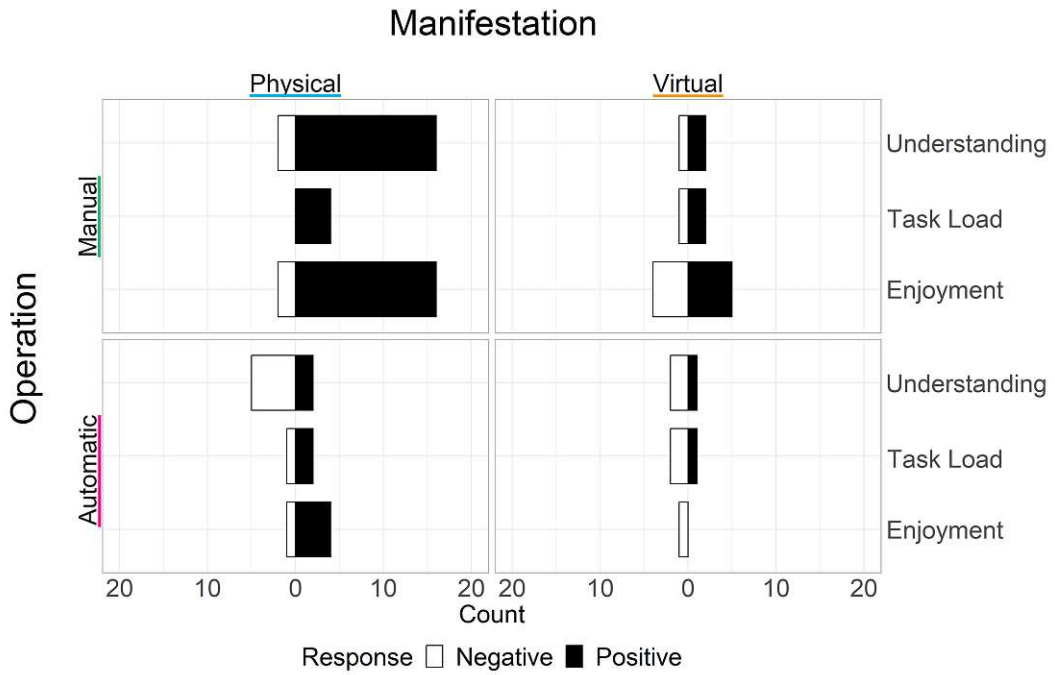

Figure 6.7: Results for the deductive coding of the qualitative feedback. The Y-axes in the subplots refer to the codes of the respective hypotheses.

negative comments respectively. Here, participants reported confusion about how automated representations would react to complications ("*frequency was missing*", "*visualization of the diastole was unclear*").

Task Load (H2). Deductive coding reveals four statements pointing towards increased task load when working with physical manual representations. The inductive coding shows three of these comments referring to increased *cognitive effort*. Examining these comments in detail, we find that only one participant specifically stated that they found added cognitive load added by the physical manual representation ("*The hands-on interaction adds an extra layer of cognitive effort*"). One participant reported difficulties in determining the difference between physical manual and automated representation. The final comment was about the thickness of the black line, and finding it hard to see the water line. For the virtual manual representation, two participants stated the desire for refined controls. The only statement about *physical effort* in the physical representation states the necessity to step back to gain a better overview.

Enjoyment (H3). 16 Participants made positive remarks on their enjoyment of the physical manual (''*Doing the work yourself was engaging*", "*Manual pump was fun*", "*[I liked the] interactivity*"), five for the virtual manual, and four for the physical automated representation. We did not receive any positive comments on the virtual automated model. Statements about *interaction* were a large group of positive remarks on both manual representations, counting 15

Table 6.1: Summary of the triangulation process.

	Quantitative	Qualitative	Triangulation
H1	Quiz	*Deductive Code:* • Understanding *Inductive Codes:* • Non-visual understanding • Visual understanding	Qualitative support
H2	NASA-TLX	*Deductive Code:* • Task load *Inductive Codes:* • Physical load • Cognitive load	Quantitative support
H3	Questionnaire, Ranking	*Deductive Code:* • Enjoyment *Inductive Codes:* • Interaction • Aesthetics	Full support

for the physical and five for the virtual. As another positive factor for the enjoyment of both of the physical models, four participants reported positively on our design, which we grouped into the category *aesthetics* ("*The clear tubes in the physical display were most compelling*").

6.4.3 Triangulation

Finally, we discuss the triangulation of our results (Table 6.1):

Understanding (H1). While quantitative methods did not show measurable improvements, the positive results for the inductive coding in the category *understanding*, as well as the positive results for physical manual representations in the inductive categories *visual understanding* and *non-visual understanding* positively support our theory. In conclusion, H1 is only supported subjectively, and we can not show a measurable knowledge gain. Thus, we recommend further research on the effect of tactile feedback in manual representations.

Task Load (H2). Here, our quantitative analysis revealed that our participants reported a higher task load when working with manually operated representations. However, our qualitative analysis does not fully support this result. Of the comments that relate to increased task load, only a single comment directly referred to a higher *cognitive effort* when working with physical manual representations in general. The remaining comments were specifically aimed at visual difficulties, resulting from the way we built our physical model. The only cognitive difficulties in working with virtual manual models, reported only by two participants, referred to difficulties with the keyboard controls. *Physical effort* did not appear as a factor in the interaction with our models. While there is some evidence to support the hypothesis that working with manual models

increases subjective task load, we see little support for this from our qualitative analysis. We therefore conclude that task load did not play an important role for our participants.

Enjoyment (H3). In terms of the subjective enjoyment of physical and manual representations, the qualitative and quantitative results align. Physical and manual representations yielded significantly higher results on our engagement questionnaire. In the summary ranking, physical manual representations were preferred by a majority. Further qualitative analysis reveals that *interaction* is a main facilitator of engagement, both in physical and virtual models. Additionally, the *aesthetics* of our representations were positively remarked upon in the physical models. The evidence collected supports the hypothesis that physical manifestation, as well as manually operated representations, are enjoyed more than virtual or automated ones.

6.5 Discussion

Our study is a first attempt to disentangle the factors of operation mode and manifestation to highlight their individual effects on representations of a complex physiological process and associated complications. While we study a very specific example derived from the medical domain, we use common physicalization concepts like an embodied metaphor [ZVM08] and direct interaction [BZW+22]. With this, we aim to make the concept understandable to a lay audience. From our results, we argue that **physical manifestation and direct manual operation are both beneficial factors in data physicalizations**.

Our quantitative metrics did not show a positive impact of physical manifestation and manual interaction on our participants' understanding of the represented process. This finding is consistent with prior studies conducted by Jansen et al. [JDF13], Hurtienne et al. [HMC+20] and Drogemuller et al. [DCW+21]. However, our qualitative methods reveal a subjective benefit, as reported by participants in their feedback. This indicates that **manual operation in the form of direct interactivity in physical representations can lead to increased understanding**. As opposed to physical representations, virtual representations do not benefit from such non-visual insight. Interestingly, our findings contrast to the results by Pollalis et al. [PMW+18], who highlight the shortcomings of physical representations in accurately representing visual information. Instead, we find that when using manual operation, virtual and physical representations can provide comparable levels of understanding. Finally, with respect to memorability, Stusak et al. [SSB15] have shown that physicalizations are more memorable than virtual visualizations. Our study did not consider such long-term effects, but **we recommend investigating further the effects of direct interactivity on the memorability of physicalizations**.

Furthermore, we show that **physical data representations benefit highly from direct interactivity**. Our quantitative results indicate higher engagement for such representations, consistent with previous findings of Hurtienne et al. [HMC+20]. Our results indicate that this is not caused by a difference in manifestation alone. **If observers are not able to manually interact with a representation, a virtual one may be equally or better suited.** This is also supported by our participants' ranking of the different representations, where virtual manual and physical automated models were ranked similarly.

In our study, we focus on metaphorical models. The creation of more complex representations would be more feasible in a purely virtual setting than in the physical space. **We think that manual interaction can serve as an enhancement to both physical and virtual representations.** Similar to Hurtienne et al. [HMC+20], who measured the physical effort to increase when working with a physical representation, we observe increased reported task load in the manual condition. On closer inspection, our quantitative results show that this may stem from the lessened use of cognitive resources of our participants when they did not interact with the representations themselves. Combined with the heightened engagement during manual interaction, we conclude that **representations without direct interaction are less engaging.**

6.6 Limitations and Future Work

We designed our representation with the assumption that modeling the cardiac cycle as a simplified metaphor would be beneficial for a layperson audience. Despite having been checked by a clinical doctor, our abstracted model and its variations may have caused misunderstandings of the (patho)physiological processes they represent. In our study, we collected 14 negative comments that referred to the model itself as opposed to aspects specific to manifestation or mode of operation. Thus, **we recommend investigating further the effect of the degree of abstraction on (physical) metaphors in future work.**

When creating our virtual representations, we modeled the pump metaphor in unity with certain concessions. While using a simple, two-dimensional model of our physical setup did not visibly impact our results, limiting the interaction with the virtual pump to the keyboard led to an absence of tactile feedback. Because of 15 comments referencing the positive impact of manual interaction on the engagement, as well as 6 comments mentioning non-visual insights gained from physical representations, a different method of controlling a virtual representation may yield more positive results. Tactile feedback may be a greater influence than immersion, augmented or virtual reality methods may offer limited benefits. **When comparing interactive representations in the future, we recommend using input devices with feedback mechanisms**. This could range from force-feedback joysticks to custom devices that simulate physical resistance.

Our analysis of the test scores failed to show an advantage of either of our four representations on the participant's understanding. We used four scenarios to test our hypotheses with four different representations. One used scenario, simulating physical exertion (**S1**), was not fully equivalent to the other three. Participants using manually operated representations scored worse in the exertion scenario. This may have been because the possibility of non-visual insight was not given because of the lack of tactile feedback. Therefore, **we recommend considering the presence of tactile feedback in manual representations in future comparative studies.**

Finally, we acknowledge that our sample was biased towards a young, age group of male subjects with a relatively high education. Cardiac pathologies are relevant to everyone especially due to preventable risk factors, such as smoking. However, data physicalization may have additional benefits in educating an audience with low visualization literacy, including children. Future work considering a **more diverse demographic** could reveal further insights into the effectiveness

of educational pathophysiological process representations—also with regard to learning aspects such as constructivism, active learning, or learning-by-doing [HJC14, HHJ+17].

Conclusion We present a study on a representation that abstracts the cardiac cycle into a tangible metaphor. Using four different versions of this representation, we investigate the effects of manifestation and mode of operation on their efficiency and engagement. Our findings show that physical, manually operated representation have the potential to engage an audience in an emotional way, through their unique way of immersing a user. This constitutes a further step towards using the intermodal potential of direct interactivity in data physicalization.

CHAPTER 7

Conclusion

In this dissertation, I explored how data physicalizations leverage *interactivity* in different ways to engage human users. My overarching research goal was to investigate the benefits and pitfalls of different ways to integrate a user's actions into data physicalizations. I presented four examples of different ways to interact with data physicalizations, progressively moving from interacting with the representation to becoming an active part of it. This section serves as a summary and discussion of the results and provides a synthesis of my findings, pointing to future research opportunities in the field.

7.1 Reflection

In Chapter 1, I posed four research questions centered around the benefits of data physicalizations with different ways of integrating user actions. I addressed each research question in distinct research papers (Chapters 3– 6), for which I summarize the outcomes here.

Q1: ***Which benefits arise from physically representing a 3D visualization for dynamic data with a modular representation?*** In Chapter 3, we designed a simple way to arrange time-sliced data in the physical space to address this question. We based our technique on existing visual metaphors such as *Time-Lighting* [FCAA23], which simulates the projection of timesliced dynamic graphs on a two-dimensional surface. Our representation instead foregoes the projection in favor of relying on natural spatial perception in the real world. Using printable overhead foils, which are available in office supply stores, makes this method accessible without the need for expensive technology. In a case study, we analyzed how our *HoloGraphs* stimulate spatial 3D perception by interviewing a domain expert in a case study. **Takeaways:** Without prior experience with dynamic graph representations, the expert was able to reason about character interactions in the Harry Potter series, pointing out how movement patterns of nodes represent certain story events. This demonstrates the ability of a physical time-sliced space-time cube to

stimulate natural spatial perception, where a visualization would have to make use of abstractions and projections.

Q2: ***What is the influence of constructing a physicalization with a pre-defined visual embedding on human perception?*** In Chapter 4, we investigated this by defining and studying a physical graph construction toolkit: **NODKANT**. Our toolkit, created by FFF 3D printing two elementary parts and using yarn and magnets, is designed to specifically support users in the construction of node-link diagrams. In our study, we compared how people construct graphs when given only data as opposed to when given a pre-defined layout. We observed that people who constructed their graphs freely would become deeply engaged in construction and actively think and reason about the evolving structures, like clusters and important nodes. Conversely, those who were given a fixed layout mostly focused on aesthetic concerns like straight edges and node positions. We also studied how these participants completed network benchmark tasks in comparison to a group that received a finished representation. **Takeaways:** Here, we saw that performance between the conditions did not differ significantly, even though freely constructed layouts were objectively messier, which may be owed to the fact that participants in these conditions interacted much more liberally with their graphs. Moreover, we observed that participants who built their graphs retained more insights about them. This indicates that tactile interaction can stimulate sensemaking, without the tactile channel being used for information encoding.

Q3: ***How does tactile perception of local elastic compliance support the exploration of three-dimensional artifacts?*** To investigate this, we designed and evaluated *squishicalization*, a fabrication pipeline for elastic volume physicalization, in Chapter 5. Using the borders of Voronoi diagrams, created from point distributions that map to a scalar field via a custom mapping function, *squishicalization* allows users to create elastic representations of volume data. Two studies on how humans perceive differences in elasticities in heterogenetically composed artifacts shed light on the potential of tactile information encoding. Users were able to accurately determine the position and size of harder regions within soft bodies and to determine the composition of nested artifacts. **Takeaways:** In a mechanical evaluation, we showed that this method can reproduce the elastic compliance of a wide variety of human tissues, pointing to a possible use for our method in anatomical props. This is further underscored by expert interviews, where we also discussed the possibility of our technique to highlight structural properties of artifacts that are not instantly visually apparent by tactile exploration. *Squishicalization* highlights how our senses can act in complementary ways when a representation makes use of direct interactivity.

Q4: ***Which are the benefits of manually operating a physical representation of a dynamic process compared to virtual and operated ones?*** Chapter 6 showed the design and evaluation of our interactive physicalization of the cardiac cycle, the *Heart Machine*. With the abstraction of the human heart as a manually or automatically operating pump, which we investigated as a physical installation as well as a virtual simulation, we conducted a between-subject study on how users understand different abnormal conditions. **Takeaways:** While we did not observe a significant advantage for physical representations in insight generation, we saw evidence that physical feedback supports understanding. With regards to engagement we showed that physical

and manual operation has a positive influence, whereas only manual operation increases task load. Users vastly preferred interacting with the physical representation, while automated and virtual ones were less positively received. This showcases that physical representations hold unique potential for the utilization of inter-modal information encoding, through active perception.

7.2 Synthesis

Physicalizations demonstrated their potential as tangible 3D visualizations in the past, for example, with physical 3D bar charts outperforming screen-based ones [JDF13]. With *HoloGraphs*, I show how a more complex concept based on space-time cubes can profit from representation in physical space. The limitations of the approach are mostly caused by the inflexibility of currently existing fabrication methods. The 3D printing devices I used in all of the presented projects are only able to produce a limited number of prototypes at a time, with printing times ranging from a few hours for a set of **NODKANT** parts, to several days for the printing of a *squishicalization*. A prototype is limited to a specific canvas size based on the available printing technology, and data cannot be added to or removed from a physical representaiton with ease. Aside from this, we managed to showcase the potential of several versatile concepts. Our interactive physicalizations serve as proof of concepts for visualization techniques that can not be implemented on currently existing hardware. However, 3D printing devices are constantly improving, making the proposed techniques more usable and accessible every day. The overhead foils in a *HoloGraph* could be replaced by transparent screens, arranged in a similar way to create the same effect on depth perception. Coupled with interactive software, the limitations of a static approach can be overcome easily. **I propose that interactive physicalizations can serve as prototypes and inspirations for interactions with potentially interesting, but expensive or currently inaccessible techniques.**

I explored a variety of possible application areas for data physicalization, ranging from data analysis to the education of both laypeople and professionals. The emotional effects of our prototypes on users depended strongly on the embodied metaphors, but I also show the part sensory integration plays in this. Zhao and Vande Moere [ZVM08] propose that physical embodiment is a means to elicit an emotional response in an audience. I examined the subjective emotional value of two different physicalizations, **NODKANT** and *Heart Machine*, measuring aspects proposed by Wang et al. [WSK+19]. *HoloGraphs* and **NODKANT** made use of an established encoding —node link diagrams— and *squishicalization* directly reproduced anatomical structures. We did not measure emotional effects in our users based on interaction with our props. Conversely, with *Heart Machine*, we showed that both physicality and interaction with the representation caused a significant increase in emotional engagement. However, with **NODKANT**, we found indications that memorability may depend on physical interaction and immersion without the influence of affective engagement. Future work should examine further the effect that different aspects of engagement have on the efficiency and memorability of data physicalizations. **I propose that integrating a user's actions into a representation can cause emotional attachment in users and that data physicalizations can be used to target emotional responses to elicit specific benefits.**

By integrating a user's actions into data physicalization, the representations I presented differ more and more from established visualizations. *HoloGraphs* mimics a 3D visualization, however, direct interactivity necessitates a further departure from the computer screen. While still based on the same concept, constructing a representation, as demonstrated by **NODKANT** is not a typical interaction afforded by classical visualizations. However, users were able to derive insights and affordances from the construction experience, which enhances interactions with the data. Similarly, *squishicalization* does not rely on typical output devices and instead relies on naturally addressing the tactile and visual sense. Here, we found that such interactions are not easily communicated by the affordances of the artifact. Moreover, the tactile sense has a complex stimulus-response relationship, which complicates data encoding on this channel. Future work is needed for new ways to improve and refine tactile encodings, but there are indications that users can adapt to such techniques successfully. **Non-visual encodings need to consider the individual strengths and weaknesses of the intended channel to facilitate better sensemaking, with future work required on maximizing the perceptibility of such mappings**

The *Heart Machine* goes beyond conventional representation and uses intermodal encoding across the visual and tactile sense. In addition to users reporting deeper engagement, they also vastly preferred the physical, manually operated version of our apparatus over others. However, our study could not highlight an advantage in user understanding. This may be owed to the fact that not all conditions in the study fully utilized the capabilities of the apparatus, with one condition not providing tactile feedback. Onboarding helps to make data visualizations be used more effectively [SGA22]; node-link diagrams, for example, have to be explained before they can be understood by a user. Physical versions of conventional representations, such as the presented network physicalizations, clearly profit from this. However, *Heart Machine* and *squishicalization* are methods that rely on interactions that humans are naturally familiar with, such as tactile exploration or manually operating a device. For example, we found that users had difficulties identifying gradient directions in continuously adapting elastic structures or that simply increasing the target frequency of a pump without tactile feedback did not relate well to physiology. This indicates that abstract visual encodings have to be learned, while physical encodings have to support the existing sensemaking capabilities of the user. **I propose that to increase the efficiency of data physicalizations with direct interactivity, it is necessary to adapt fabrication pipelines to the sensory capabilities of humans through more complex data transformations.**

The physicalization prototypes presented in this dissertation embody data from different contexts and with different levels of abstraction. On the one hand, *HoloGraphs* and **NODKANT** deal with graphs, i.e., **abstract** representations of relationships. On the other hand, *squishicalization* and the *Heart Machine* are representations of medical data. Both representations act as physical **metaphors** for the data — *squishicalization* represents volume data in a spatial context, and the *Heart Machine* represents the work of the heart using a pump. Graphs are abstract, learned representations that have no corresponding real-world interactions, i.e., a node does not appear naturally in the physical world. Indirect interaction is more common with such representations, and encoded data is clearly communicated. Using physical metaphors such as *Heart Machine*, have inherent physical affordances, i.e., the body can be touched and a pump can be operated.

However, the encoded information may not be understood as intuitively. This points towards possible combinations in future work, e.g., Node attributes could be encoded using metaphors for the real world context in an abstract graph. **I propose that combining abstract traditional representations and metaphors representing context in data physicalization warrants future exploration.**

7.3 Closing Remarks

Passive and purely visual physicalizations like *HoloGraphs* show a first step into affording interactions with physical artifacts, while more innovative and novel contributions like *squishicalization* or the *Heart Machine* rely on integrating the user's action directly into the representation. Both sides of this spectrum can bring unique benefits to the sensemaking process. Even purely visual representations can stimulate depth perception and motivate innovative interactions, but they often demonstrate technical limitations that hold back their efficiency. On the other hand, creating immersive representations can stimulate human senses more naturally, but utilizing non-visual senses, such as tactile perception requires careful understanding of their capabilities. My dissertation shows that data physicalizations have a vast potential to serve as sense-making tools, supported by novel ways of considering user interaction. Supported by technological improvements in the future and with careful consideration of human capabilities, data physicalization is a catalyst for sensemaking to transcend beyond the screens.

Overview of Generative AI Tools Used

I have used Grammarly (`https://app.grammarly.com/`) for grammar and spelling checks in the course of writing this dissertation.

List of Figures

List of Tables

Bibliography

[AMA22] Alessio Arleo, Silvia Miksch, and Daniel Archambault. Event-based dynamic graph drawing without the agonizing pain. *Computer Graphics Forum*, 41(6):226–244, 2022.

[ANMMG25] Nora Al-Naami, Nicolas Médoc, Matteo Magnani, and Mohammad Ghoniem. Improved visual saliency of graph clusters with orderable node-link layouts. *IEEE Transactions on Visualization and Computer Graphics*, 31(1):1028–1038, 2025.

[AP13] Daniel Archambault and Helen C. Purchase. Mental map preservation helps user orientation in dynamic graphs. In Walter Didimo and Maurizio Patrignani, editors, *Graph Drawing*, pages 475–486, Berlin, Heidelberg, 2013.

[APP11] Daniel Archambault, Helen C. Purchase, and Bruno Pinaud. Animation, small multiples, and the effect of mental map preservation in dynamic graphs. *IEEE Transactions on Visualization and Computer Graphics*, 17(4):539–552, 2011.

[APS14] Jae-wook Ahn, Catherine Plaisant, and Ben Shneiderman. A task taxonomy for network evolution analysis. *IEEE Transactions on Visualization and Computer Graphics*, 20(3):365–376, 2014.

[ASS+19] Kathleen D. Ang, Faramarz F. Samavati, Samin Sabokrohiyeh, Julio Garcia, and Mohammed S. Elbaz. Physicalizing cardiac blood flow data via 3d printing. *Computers & Graphics*, 85:42–54, 2019.

[ASSB+23] Mashael AlKadi, Vanessa Serrano, James Scott-Brown, Catherine Plaisant, Jean-Daniel Fekete, Uta Hinrichs, and Benjamin Bach. Understanding barriers to network exploration with visualization: A report from the trenches. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):907–917, 2023.

[BBDW17] Fabian Beck, Michael Burch, Stephan Diehl, and Daniel Weiskopf. A taxonomy and survey of dynamic graph visualization. *Comput. Graph. Forum*, 36(1):133–159, January 2017.

[BBO+10] Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. Design and fabrication of materials with desired deformation behavior. *ACM Trans. Graph.*, 29(4), July 2010.

[BCB07] Stephen Brewster, Faraz Chohan, and Lorna Brown. Tactile feedback for mobile interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, page 159–162, New York, NY, USA, 2007.

[BDA+17] Benjamin Bach, Pierre Dragicevic, Daniel Archambault, Christophe Hurter, and Sheelagh Carpendale. A descriptive framework for temporal data visualizations based on generalized space-time cubes. *Computer Graphics Forum*, 36(6):36–61, 2017.

[BFY+24] S. Sandra Bae, Takanori Fujiwara, Anders Ynnerman, Ellen Yi-Luen Do, Michael L. Rivera, and Danielle Albers Szafir. A computational design pipeline to fabricate sensing network physicalizations. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):913–923, 2024.

[BIM12] Ulrik Brandes, Natalie Indlekofer, and Martin Mader. Visualization methods for longitudinal social networks and stochastic actor-oriented modeling. *Social Networks*, 34(3):291–308, 2012.

[BM13] Matthew Brehmer and Tamara Munzner. A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2376–2385, 2013.

[BMBH16] Katy Börner, Adam Maltese, Russell Nelson Balliet, and Joe Heimlich. Investigating aspects of data visualization literacy using 20 information visualizations and 273 science museum visitors. *Information Visualization*, 15(3):198–213, 2016.

[BOH11] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. D³ data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011.

[BPF14] Benjamin Bach, Emmanuel Pietriga, and Jean-Daniel Fekete. Graphdiaries: Animated transitions and temporal navigation for dynamic networks. *IEEE Transactions on Visualization and Computer Graphics*, 20(5):740–754, 2014.

[Bri07] Robert Bridson. Fast poisson disk sampling in arbitrary dimensions. In *ACM SIGGRAPH 2007 Sketches*, SIGGRAPH '07, page 22–es, New York, NY, USA, 2007.

[BS09] Jean-Michel Boucheix and Emmanuel Schneider. Static and animated presentations in learning dynamic mechanical systems. *Learning and Instruction*, 19(2):112–127, April 2009.

[BSP19] Emilio Barchiesi, Mario Spagnuolo, and Luca Placidi. Mechanical metamaterials: a state of the art. *Mathematics and Mechanics of Solids*, 24(1):212–234, January 2019.

[BVKVH24] Hans Brombacher, Rosa Van Koningsbruggen, Steven Vos, and Steven Houben. SensorBricks: a collaborative tangible sensor toolkit to support the development of data literacy. In *Proceedings of the Eighteenth International Conference on Tangible, Embedded, and Embodied Interaction*, TEI '24, New York, NY, USA, 2024.

[BZP+20] Fearn Bishop, Johannes Zagermann, Ulrike Pfeil, Gemma Sanderson, Harald Reiterer, and Uta Hinrichs. Construct-A-Vis: Exploring the free-form visualization processes of children. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):451–460, 2020.

[BZW+22] S. Sandra Bae, Clement Zheng, Mary Etta West, Ellen Yi-Luen Do, Samuel Huron, and Danielle Albers Szafir. Making data tangible: A cross-disciplinary design space for data physicalization. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22, New York, NY, USA, 2022.

[Chi97] Michelene T.H. Chi. Quantifying qualitative analyses of verbal data: A practical guide. *Journal of the Learning Sciences*, 6(3):271–315, 1997.

[CMS99] Stuart Card, Jock Mackinlay, and Ben Shneiderman. *Readings in Information Visualization: Using Vision To Think*. Morgan Kaufmann, 1999.

[DCW+21] Adam Drogemuller, Andrew Cunningham, James A Walsh, James Baumeister, Ross T. Smith, and Bruce H Thomas. Haptic and visual comprehension of a 2d graph layout through physicalisation. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA, 2021.

[DCWT19] Adam Drogemuller, Andrew Cunningham, James Walsh, and Bruce Thomas. Towards instantiating design principles for physical networks. In *Proceedings of the CHI 2019 Workshop on Interaction Design and Prototyping for Immersive Analytics*, 2019.

[DGL+97] Giuseppe Di Battista, Ashim Garg, Giuseppe Liotta, Roberto Tamassia, Emanuele Tassinari, and Francesco Vargiu. An experimental comparison of four graph drawing algorithms. *Computational Geometry*, 7(5):303–325, 1997.

[DJ12] Pierre Dragicevic and Yvonne Jansen. List of physical visualizations. `http://dataphys.org/list/`, 2012. Accessed: 2025-07-18.

[DJVM20] Pierre Dragicevic, Yvonne Jansen, and Andrew Vande Moere. Data physicalization. In Jean Vanderdonckt, Philippe Palanque, and Marco Winckler, editors, *Handbook of Human Computer Interaction*, pages 1–51. Springer International Publishing, Cham, 2020.

[DLF+09] Tim Dwyer, Bongshin Lee, Danyel Fisher, Kori Inkpen Quinn, Petra Isenberg, George Robertson, and Chris North. A comparison of user-generated and automatic graph layouts. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):961–968, 2009.

[DLM24] Walter Didimo, Giuseppe Liotta, and Fabrizio Montecchiani. Chapter 1: Network data visualization. In *Handbook of Social Computing*, pages 2 – 11. Edward Elgar Publishing, Cheltenham, UK, 2024.

[DMSJ23] Hessam Djavaherpour, Lynn Moorman, Faramarz Samavati, and Yvonne Jansen. First insights into intuit: An interactive tactile physicalization for user interpretation of radar technology. *IEEE Computer Graphics and Applications*, 43(5):91–98, 2023.

[DSMA+21] Hessam Djavaherpour, Faramarz Samavati, Ali Mahdavi-Amiri, Fatemeh Yazdanbakhsh, Samuel Huron, Richard Levy, Yvonne Jansen, and Lora Oehlberg. Data to physicalization: A survey of the physical rendering process. *Computer Graphics Forum*, 40(3):569–598, 2021.

[EBK+25] Henry Ehlers, Nicolas Brich, Michael Krone, Martin Nöllenburg, Jiacheng Yu, Hiroaki Natsukawa, Xiaoru Yuan, and Hsiang-Yun Wu. An introduction to and survey of biological network visualization. *Computers & Graphics*, 126:104115, 2025.

[EEC+24] Kajetan Enge, Elias Elmquist, Valentina Caiola, Niklas Rönnberg, Alexander Rind, Michael Iber, Sara Lenzi, Fangfei Lan, Robert Höldrich, and Wolfgang Aigner. Open your ears and take a look: A state-of-the-art report on the integration of sonification and visualization. *Computer Graphics Forum*, 43(3):e15114, 2024.

[EMWR24] Henry Ehlers, Diana Marin, Hsiang-Yun Wu, and Renata Raidou. Visualizing group structure in compound graphs: The current state, lessons learned, and outstanding opportunities. In *Proceedings of the 19th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - IVAPP*, pages 697–708, 2024.

[FABM24] Velitchko Filipov, Alessio Arleo, Markus Bögl, and Silvia Miksch. On network structural and temporal encodings: A space and time odyssey. *IEEE Transactions on Visualization and Computer Graphics*, 30(8):5847–5860, 2024.

[FAM23] Velitchko Filipov, Alessio Arleo, and Silvia Miksch. Are we there yet? a roadmap of network visualization from surveys to task taxonomies. *Computer Graphics Forum*, 42(6):e14794, 2023.

[FCAA23] Velitchko Filipov, Davide Ceneda, Daniel Archambault, and Alessio Arleo. Timelighting: Guidance-enhanced exploration of 2d projections of temporal graphs. In Michael A. Bekos and Markus Chimani, editors, *Graph Drawing and Network Visualization*, pages 231–245, Cham, 2023.

[FH95] Erika L. Ferguson and Mary Hegarty. Learning with real machines or diagrams: Application of knowledge to real-world problems. *Cognition and Instruction*, 13(1):129–160, 1995.

[Fox06] Robert Fox. The history of science, medicine and technology at oxford. *Notes and records of the Royal Society of London*, 60:69–83, 02 2006.

[Fri37] Milton Friedman. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association*, 32(200):675–701, 1937.

[Fri22] Sarah F. Frisken. Surfacenets for multi-label segmentations with preservation of sharp boundaries. *The Journal of Computer Graphics Techniques*, 11(1):34–54, 2022.

[Gau14] David Gauntlett. The LEGO System as a tool for thinking, creativity, and changing the world. In *Lego studies: Examining the building blocks of a transmedial phenomenon*, pages 189–205, New York, 2014.

[GCC+10] Elizabeth Gire, Adrian Carmichael, Jackie Chini, Amy Rouinfar, N. Sanjay Rebello, Garrett Smith, and Sadhana Puntambekar. The effects of physical and virtual manipulatives on students' conceptual learning about pulleys. In *Proceedings of the 9th International Conference of the Learning Sciences (ICLS 2010)*, volume 1, pages 937–943, June 2010.

[GGMR20] Carlos F. Guimarães, Luca Gasperini, Alexandra P. Marques, and Rui L. Reis. The stiffness of living tissues and its implications for tissue engineering. *Nature Reviews Materials*, 5(5):351–370, May 2020.

[GTS10] Lars Grammel, Melanie Tory, and Margaret-Anne Storey. How information visualization novices construct visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):943–952, 2010.

[GV16] Julia Anne Garde and Mascha Cecile van der Voort. Could LEGO® Serious Play® be a useful technique for product co - design? *DRS Biennial Conference Series*, 2016.

[GYM+17] Justine Garcia, ZhiLin Yang, Rosaire Mongrain, Richard L Leask, and Kevin Lachapelle. 3D printing materials and their use in medical education: a review of current technology and trends for the future. *BMJ Simulation & Technology Enhanced Learning*, 4(1):27–40, December 2017.

[Han13] Insook Han. Embodiment: A new perspective for evaluating physicality in learning. *Journal of Educational Computing Research*, 49(1):41–59, 2013.

[Har03] Rex Hartson. Cognitive, physical, sensory, and functional affordances in interaction design. *Behaviour & Information Technology*, 22(5):315–338, 2003.

[HCBF16] Samuel Huron, Sheelagh Carpendale, Jeremy Boy, and Jean-Daniel Fekete. Using viskit: A manual for running a constructive visualization workshop. In *Pedagogy of Data Visualization Workshop at IEEE VIS 2016*, Baltimore, MD, United States, October 2016.

[HCT+14] Samuel Huron, Sheelagh Carpendale, Alice Thudt, Anthony Tang, and Michael Mauerer. Constructive visualization. In *Proceedings of the 2014 conference on Designing interactive systems*, DIS '14, pages 433–442, New York, NY, USA, June 2014.

[HGG+16] Steven Houben, Connie Golsteijn, Sarah Gallacher, Rose Johnson, Saskia Bakker, Nicolai Marquardt, Licia Capra, and Yvonne Rogers. Physikit: Data engagement through physical ambient visualizations in the home. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, page 1608–1619, New York, NY, USA, 2016. Association for Computing Machinery.

[HH16] Trevor Hogan and Eva Hornecker. Towards a design space for multisensory data representation. *Interacting with Computers*, 29(2):147–167, 05 2016.

[HHJ+17] Trevor Hogan, Uta Hinrichs, Yvonne Jansen, Samuel Huron, Pauline Gourlet, Eva Hornecker, and Bettina Nissen. Pedagogy & physicalization: Designing learning activities around physical data representations. In *Proceedings of the 2017 ACM Conference Companion Publication on Designing Interactive Systems*, DIS '17 Companion, page 345–347, New York, NY, USA, 2017.

[HJC14] Samuel Huron, Yvonne Jansen, and Sheelagh Carpendale. Constructing visual representations: Investigating the use of tangible tokens. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2102–2111, 2014.

[HL22] Said Hajouli and Dipesh Ludhwani. Heart failure and ejection fraction. In *StatPearls*. Treasure Island (FL), 2022.

[HMC+20] Jörn Hurtienne, Franzisca Maas, Astrid Carolus, Daniel Reinhardt, Cordula Baur, and Carolin Wienrich. Move&find: The value of kinaesthetic experience in a casual data representation. *IEEE Computer Graphics and Applications*, 40(6):61–75, 2020.

[HMHU13] Aki Hayashi, Tatsushi Matsubayashi, Takahide Hoshide, and Tadasu Uchiyama. Initial positioning method for online and real-time dynamic graph drawing of time varying data. In *2013 17th International Conference on Information Visualisation*, pages 435–444, 2013.

[HS88] Sandra G. Hart and Lowell E. Staveland. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Advances in Psychology*, volume 52 of *Human Mental Workload*, pages 139–183. January 1988.

[IKS+18] Alexandra Ion, Robert Kovacs, Oliver S. Schneider, Pedro Lopes, and Patrick Baudisch. Metamaterial textures. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, page 1–12, New York, NY, USA, 2018.

[JD13] Yvonne Jansen and Pierre Dragicevic. An interaction model for visualizations beyond the desktop. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2396–2405, 2013.

[JDF13] Yvonne Jansen, Pierre Dragicevic, and Jean-Daniel Fekete. Evaluating the efficiency of physical visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, page 2593–2602, New York, NY, USA, 2013.

[JDI+15] Yvonne Jansen, Pierre Dragicevic, Petra Isenberg, Jason Alexander, Abhijit Karnik, Johan Kildal, Sriram Subramanian, and Kasper Hornbæk. Opportunities and challenges for data physicalization. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, page 3227–3236, New York, NY, USA, 2015.

[KAL+16] Rohit Ashok Khot, Josh Andres, Jennifer Lai, Juerg von Kaenel, and Florian 'Floyd' Mueller. Fantibles: Capturing cricket fan's story in 3d. In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, DIS '16, page 883–894, New York, NY, USA, 2016.

[KET+21] Hyunyoung Kim, Aluna Everitt, Carlos Tejada, Mengyu Zhong, and Daniel Ashbrook. Morpheesplug: A toolkit for prototyping shape-changing interfaces. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA, 2021. Association for Computing Machinery.

[KHM14] Rohit Ashok Khot, Larissa Hjorth, and Florian 'Floyd' Mueller. Understanding physical activity through 3d printed material artifacts. page 3835–3844, 2014.

[Kir10] David Kirsh. Thinking with external representations. *AI & SOCIETY*, 25(4):441–454, November 2010.

[KL22] Jordan King and David R. Lowery. Physiology, Cardiac Output. In *StatPearls*. Treasure Island (FL), 2022.

[Kri18] Klaus Krippendorff. Content analysis: An introduction to its methodology. 2018.

[KSB+23] Magdaléna Kejstová, Christina Stoiber, Magdalena Boucher, Martin Kandlhofer, Simone Kriglstein, and Wolfgang Aigner. Construct and play: Engaging students with visualizations through playful methods. In *Companion Proceedings of the Annual Symposium on Computer-Human Interaction in Play*, CHI PLAY Companion '23, page 96–101, New York, NY, USA, 2023. Association for Computing Machinery.

[KWK21] Maria Karyda, Danielle Wilde, and Mette Gislev Kjærsgaard. Narrative physicalization: Supporting interactive engagement with personal data. *IEEE Computer Graphics and Applications*, 41(1):74–86, 2021.

[LC98] William E. Lorensen and Harvey E. Cline. Marching cubes: a high resolution 3D surface construction algorithm. In *Seminal graphics*, pages 347–353. New York, NY, USA, July 1998.

[LKG+16] Patric Ljung, Jens Krüger, Eduard Groller, Markus Hadwiger, Charles D. Hansen, and Anders Ynnerman. State of the art in transfer functions for direct volume rendering. *Computer Graphics Forum*, 35(3):669–691, 2016.

[LPP+06] Bongshin Lee, Catherine Plaisant, Cynthia Sims Parr, Jean-Daniel Fekete, and Nathalie Henry. Task taxonomy for graph visualization. In *Proceedings of the 2006 AVI Workshop on BEyond Time and Errors: Novel Evaluation Methods for Information Visualization*, BELIV '06, page 1–5, New York, NY, USA, 2006.

[LSZ+14] Lin Lu, Andrei Sharf, Haisen Zhao, Yuan Wei, Qingnan Fan, Xuelin Chen, Yann Savoye, Changhe Tu, Daniel Cohen-Or, and Baoquan Chen. Build-to-last: strength to weight 3D printed objects. *ACM Trans. Graph.*, 33(4):97:1–97:10, July 2014.

[Max95] Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995.

[McA16] Britni R. McAshan. The health museum debuts 12-foot giant heart to kick off heart health month. `https://www.tmc.edu/news/2016/01/the-health-museum-debuts-12-foot-giant-heart-to-kick-off-heart-health-month/`, January 2016. Accessed: 2025-22-07.

[MDL16] Jonàs Martínez, Jérémie Dumas, and Sylvain Lefebvre. Procedural voronoi foams for additive manufacturing. *ACM Transactions on Graphics*, 35(4):1–12, July 2016.

[Mey79] Wiliam James Meyers. Topological sorting and well-formed strings1. *Fundamenta Informaticae*, 2(1):199–209, 1979.

[MFD20] Filippo Menczer, Santo Fortunato, and Clayton A. Davis. *A First Course in Network Science*. 1 edition, January 2020.

[MGS+21] Monique Meuschke, Laura Garrison, Noeska Smit, Stefan Bruckner, Kai Lawonn, and Bernhard Preim. Towards narrative medical visualization. *arXiv preprint arXiv:2108.05462*, 2021.

[MGWP23] Jacob Miller, Mohammad Ghoniem, Hsiang-Yun Wu, and Helen C. Purchase. On the perception of small sub-graphs. In Michael A. Bekos and Markus Chimani, editors, *Graph Drawing and Network Visualization*, pages 213–230, Cham, 2023.

[MHSL18] Jonàs Martínez, Samuel Hornus, Haichuan Song, and Sylvain Lefebvre. Polyhedral voronoi diagrams for additive manufacturing. *ACM Trans. Graph.*, 37(4), July 2018.

[MNA12] Daniel Mochon, Michael I. Norton, and Dan Ariely. Bolstering and restoring feelings of competence via the IKEA effect. *International Journal of Research in Marketing*, 29(4):363–369, December 2012.

[Moe08] Andrew Vande Moere. Beyond the tyranny of the pixel: Exploring the physicality of information visualization. In *2008 12th International Conference Information Visualisation*, pages 469–474, 2008.

[Moo16] WH Moolman. The graeco-latin square and hyper graeco-latin square designs. *Journal of Mathematics and Statistical Science*, 6(8):211–220, 2016.

[MSF24] Michael J. McGuffin, Ryan Servera, and Marie Forest. Path tracing in 2d, 3d, and physicalized networks. *IEEE Transactions on Visualization and Computer Graphics*, 30(7):3564–3577, 2024.

[NEFM99] Rafael E. Núñez, Laurie D. Edwards, and João Filipe Matos. Embodied cognition as grounding for situatedness and context in mathematics education. *Educational Studies in Mathematics*, 39(1):45–65, June 1999.

[NMJ+15] Leyla Norooz, Matthew Louis Mauriello, Anita Jorgensen, Brenna McNally, and Jon E. Froehlich. Bodyvis: A new approach to body learning through wearable sensing and visualization. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, page 1025–1034, New York, NY, USA, 2015. Association for Computing Machinery.

[Nor06] C. North. Toward measuring visualization insight. *IEEE Computer Graphics and Applications*, 26(3):6–9, 2006.

[NP16] Deborah Nolan and Jamis Perrett. Teaching and learning data visualization: Ideas and assignments. *The American Statistician*, 70(3):260–269, 2016.

[Off20] Dietmar Offenhuber. What we talk about when we talk about data physicality. *IEEE Computer Graphics and Applications*, 40(6):25–37, 2020.

[OKK13] D. Oelke, D. Kokkinakis, and D. A. Keim. Fingerprint matrices: Uncovering the dynamics of social networks in prose literature. *Computer Graphics Forum*, 32(3pt4):371–380, 2013.

[OMN10] Alicia O'Cathain, Elizabeth Murphy, and Jon Nicholl. Three techniques for integrating data in mixed methods studies. *Bmj*, 341, 2010.

[OSF04] Claire O'Malley and Danae Stanton Fraser. Literature Review in Learning with Tangible Technologies, 2004. A NESTA Futurelab Research report - report 12.

[PA22] Sai Harika Pujari and Pradyumna Agasthi. Aortic Stenosis. In *StatPearls*. Treasure Island (FL), 2022.

[Pap93] Seymour Papert. *Mindstorms: Children, Computers, And Powerful Ideas*. 1993.

[PEW+24] Daniel Pahr, Henry Ehlers, Hsiang-Yun Wu, Manuela Waldner, and RenataG. Raidou. Investigating the effect of operation mode and manifestation on physicalizations of dynamic processes. *Computer Graphics Forum*, 43(3):e15106, 2024.

[PLR+16] Michal Piovarči, David I. W. Levin, Jason Rebello, Desai Chen, Roman Ďurikovič, Hanspeter Pfister, Wojciech Matusik, and Piotr Didyk. An interaction-aware, perceptual model for non-linear elastic objects. *ACM Transactions on Graphics*, 35(4):1–13, July 2016.

[PMW+18] Christina Pollalis, Elizabeth Joanna Minor, Lauren Westendorf, Whitney Fahnbulleh, Isabella Virgilio, Andrew L. Kun, and Orit Shaer. Evaluating learning with tangible and virtual representations of archaeological artifacts. In *Proceedings of the Twelfth International Conference on Tangible, Embedded, and Embodied Interaction*, TEI '18, page 626–637, New York, NY, USA, 2018. Association for Computing Machinery.

[PPP12] Helen C. Purchase, Christopher Pilcher, and Beryl Plimmer. Graph drawing aesthetics—created by users, not algorithms. *IEEE Transactions on Visualization and Computer Graphics*, 18(1):81–92, 2012.

[PRC+23] Laura J Perovich, Bernice Rogowitz, Victoria Crabb, Jack Vogelsang, Sara Hartleben, and Dietmar Offenhuber. The tactile dimension: a method for physicalizing touch behaviors. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, pages 1–15, New York, NY, USA, April 2023.

[Pun02] Samantha Punch. Research with children: The same or different from research with adults? *Childhood*, 9(3):321–341, 2002.

[Pur14] Helen Purchase. A healthy critical attitude: Revisiting the results of a graph drawing study. *Journal of Graph Algorithms and Applications*, 18(2):281–311, 2014.

[PWB21] Laura J. Perovich, Sara Ann Wylie, and Roseann Bongiovanni. Chemicals in the Creek: designing a situated data physicalization of open government data with the community. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):913–923, February 2021.

[PWR21] Daniel Pahr, Hsiang-Yun Wu, and Renata Georgia Raidou. Vologram: An educational holographic sculpture for volumetric medical data physicalization. In *VCBM 2021: 11th Eurographics Workshop on Visual Computing for Biology and Medicine, Paris, France, 22-24 September 2021*, pages 19–23, 2021.

[RA15a] Ryan Rossi and Nesreen Ahmed. The network data repository with interactive graph analytics and visualization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 29(1), Mar. 2015.

[RA15b] Ryan A. Rossi and Nesreen K. Ahmed. The Network Data Repository with Interactive Graph Analytics and Visualization. http://web.archive.org/web/20080207010024/http://www.808multimedia.com/winnt/kernel.htm, 2015. Accessed: 2025-07-22.

[Rau20] Martina A Rau. Comparing multiple theories about learning with physical and virtual representations: Conflicting or complementary effects? *Educational Psychology Review*, 32(2):297–325, 2020.

[RGW20] Renata G. Raidou, M. Eduard Gröller, and Hsiang-Yun Wu. Slice and dice: A physicalization workflow for anatomical edutainment. *Computer Graphics Forum*, 39(7):623–634, 2020.

[RHGC15] Jennifer J. H. Reynolds, Ben T. Hirsch, Stanley D. Gehrt, and Meggan E. Craft. Raccoon contact networks predict seasonal susceptibility to rabies outbreaks and limitations of vaccination. *Journal of Animal Ecology*, 84(6):1720–1731, 2015.

[RM14] Sébastien Rufiange and Guy Melançon. Animatrix: A matrix-based visualization of software evolution. In *2014 Second IEEE Working Conference on Software Visualization*, pages 137–146, 2014.

[Rot20] Gregory A. et al. Roth. Global Burden of Cardiovascular Diseases and Risk Factors, 1990–2019. *Journal of the American College of Cardiology*, 76(25):2982–3021, December 2020.

[RUK+10] Markus Rohrschneider, Alexander Ullrich, Andreas Kerren, Peter F. Stadler, and Gerik Scheuermann. Visual network analysis of dynamic metabolic pathways. In George Bebis, Richard Boyle, Bahram Parvin, Darko Koracin, Ronald Chung, Riad Hammoud, Muhammad Hussain, Tan Kar-Han, Roger Crawfis, Daniel Thalmann, David Kao, and Lisa Avila, editors, *Advances in Visual Computing*, pages 316–327, Berlin, Heidelberg, 2010.

[SAK20] Paolo Simonetto, Daniel Archambault, and Stephen Kobourov. Event-based dynamic graph visualisation. *IEEE Transactions on Visualization and Computer Graphics*, 26(7):2373–2386, 2020.

[SARW+15] Pitchaya Sitthi-Amorn, Javier E. Ramos, Yuwang Wangy, Joyce Kwan, Justin Lan, Wenshou Wang, and Wojciech Matusik. MultiFab: a machine vision assisted platform for multi-material 3D printing. *ACM Transactions on Graphics*, 34(4):1–11, July 2015.

[SB17] Sergej Stoppel and Stefan Bruckner. Vol2velle: Printable interactive volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):861–870, 2017.

[SBR+15] Christian Schumacher, Bernd Bickel, Jan Rys, Steve Marschner, Chiara Daraio, and Markus Gross. Microstructures to control elasticity in 3D printing. *ACM Transactions on Graphics*, 34(4):136:1–136:13, July 2015.

[SCP+16] Hiroki Sayama, Catherine Cramer, Mason A. Porter, Lori Sheetz, and Stephen Uzzo. What are essential concepts about networks? *Journal of Complex Networks*, 4(3):457–474, September 2016.

[SFJ+22] Stacey D. Schulberg, Amy V. Ferry, Kai Jin, Lucy Marshall, Lis Neubeck, Fiona E. Strachan, and Nicholas L. Mills. Cardiovascular risk communication strategies in primary prevention. A systematic review with narrative synthesis. *Journal of Advanced Nursing*, 78(10):3116–3140, October 2022.

[SGA22] Christina Stoiber, Florian Grassinger, and Wolfgang Aigner. Abstract and concrete materials: What to use for visualization onboarding for a treemap visualization? In *Proceedings of the 15th International Symposium on Visual Information Communication and Interaction*, VINCI '22, New York, NY, USA, 2022. Association for Computing Machinery.

[SGBI22] Mickael Sereno, Stéphane Gosset, Lonni Besançon, and Tobias Isenberg. Hybrid touch/tangible spatial selection in augmented reality. *Computer Graphics Forum*, 41(3):403–415, 2022.

[Shn03] Ben Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In Benjamin B. Bederson and Ben Shneiderman, editors, *The Craft of Information Visualization*, Interactive Technologies, pages 364–371. San Francisco, 2003.

[SK19] C. Sullivan and Alexander Kaszynski. PyVista: 3D plotting and mesh analysis through a streamlined interface for the Visualization Toolkit (VTK). *Journal of Open Source Software*, 4(37):1450, May 2019.

[SKA23] Miriam Sturdee, Hayat Kara, and Jason Alexander. Exploring co-located interactions with a shape-changing bar chart. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, New York, NY, USA, 2023. Association for Computing Machinery.

[SKRW22] Marwin Schindler, Torsten Korpitsch, Renata G. Raidou, and Hsiang-Yun Wu. Nested papercrafts for anatomical and biological edutainment. *Computer Graphics Forum*, 41(3):541–553, 2022.

[SSB15] Simon Stusak, Jeannette Schwarz, and Andreas Butz. Evaluating the memorability of physical visualizations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, page 3247–3250, New York, NY, USA, 2015.

[SSJ+14] Saiganesh Swaminathan, Conglei Shi, Yvonne Jansen, Pierre Dragicevic, Lora Oehlberg, and Jean-Daniel Fekete. Creating physical visualizations with makervis. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '14, pages 543–546, New York, NY, USA, April 2014.

[ST02] Masaki Suwa and Barbara Tversky. External representations contribute to the dynamic construction of ideas. In Mary Hegarty, Bernd Meyer, and N. Hari Narayanan, editors, *Diagrammatic Representation and Inference*, pages 341–343, Berlin, Heidelberg, 2002.

[Sta14] John Stasko. Value-driven evaluation of visualizations. In *Proceedings of the Fifth Workshop on Beyond Time and Errors: Novel Evaluation Methods for Visualization*, BELIV '14, pages 46–53, New York, NY, USA, November 2014.

[STS+14] Simon Stusak, Aurélien Tabard, Franziska Sauka, Rohit Ashok Khot, and Andreas Butz. Activity sculptures: Exploring the impact of physical visualizations on running activity. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2201–2210, 2014.

[SWR20] Marwin Schindler, Hsiang-Yun Wu, and Renata G. Raidou. The anatomical edutainer. In *2020 IEEE Visualization Conference (VIS)*, pages 1–5, 2020.

[SWW+15] Lei Shi, Chen Wang, Zhen Wen, Huamin Qu, Chuang Lin, and Qi Liao. 1.5d egocentric dynamic network visualization. *IEEE Transactions on Visualization and Computer Graphics*, 21(5):624–637, 2015.

[Tam13] Roberto Tamassia. *Handbook of graph drawing and visualization*. CRC press, 2013.

[TCKP15] Cesar Torres, Tim Campbell, Neil Kumar, and Eric Paulos. Hapticprint: Designing feel aesthetics for digital fabrication. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, UIST '15, page 583–591, New York, NY, USA, 2015. Association for Computing Machinery.

[Tec] Unity Technologies. Unity real-time development platform. https://unity.com. Accessed: 2023-03-24.

[TJW+17] Faisal Taher, Yvonne Jansen, Jonathan Woodruff, John Hardy, Kasper Hornbæk, and Jason Alexander. 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, 2017.

[TRBA20] Carlos E. Tejada, Raf Ramakers, Sebastian Boring, and Daniel Ashbrook. Air-touch: 3d-printed touch-sensitive objects using pneumatic sensing. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, page 1–10, New York, NY, USA, 2020.

[TTZ+20] Thibault Tricard, Vincent Tavernier, Cédric Zanni, Jonàs Martínez, Pierre-Alexandre Hugron, Fabrice Neyret, and Sylvain Lefebvre. Freely orientable microstructures for designing deformable 3D prints. *ACM Transactions on Graphics*, 39(6):1–16, 2020.

[VBSW13] Corinna Vehlow, Michael Burch, Hansjorg Schmauder, and Daniel Weiskopf. Radial layered matrix visualization of dynamic graphs. In *2013 17th International Conference on Information Visualisation*, pages 51–58, 2013.

[VCI20] Ivan Viola, Min Chen, and Tobias Isenberg. Visual abstraction. In Min Chen, Helwig Hauser, Penny Rheingans, and Gerik Scheuermann, editors, *Foundations of Data Visualization*, pages 15–37. Cham, 2020.

[vHR08] Frank van Ham and Bernice Rogowitz. Perceptual organization in user-generated graph layouts. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1333–1339, 2008.

[VMP10] Andrew Vande Moere and Stephanie Patel. The physical visualization of information: Designing data sculptures in an educational context. In Mao Lin Huang, Quang Vinh Nguyen, and Kang Zhang, editors, *Visual Information Communication*, pages 1–23, Boston, MA, 2010.

[VQL+06] Norbert F. Voelkel, Robert A. Quaife, Leslie A. Leinwand, Robyn J. Barst, Michael D. McGoon, Daniel R. Meldrum, Jocelyn Dupuis, Carlin S. Long, Lewis J. Rubin, Frank W. Smart, Yuichiro J. Suzuki, Mark Gladwin, Elizabeth M. Denholm, and Dorothy B. Gail. Right ventricular function and failure. *Circulation*, 114(17):1883–1891, 2006.

[WBH24] Jelle Wijers, Hans Brombacher, and Steven Houben. DataChest: a constructive data physicalization toolkit. In *Proceedings of the Eighteenth International Conference on Tangible, Embedded, and Embodied Interaction*, TEI '24, pages 1–7, New York, NY, USA, February 2024.

[WFGH11] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 143–146, Vancouver BC Canada, May 2011.

[WG10] Anne Waugh and Allison Grant. *Ross & Wilson Anatomy and physiology in health and illness*. Elsevier Health Sciences, 2010.

[WGEJ21] Michael Wüthrich, Maurus Gubser, Wilfried J. Elspass, and Christian Jaeger. A novel slicing strategy to print overhangs without support material. *Applied Sciences*, 11(18), 2021.

[WH16] Wesley Willett and Samuel Huron. A constructive classroom exercise for teaching infovis. In *Pedagogy of Data Visualization Workshop at IEEE VIS 2016*, Pedagogy of Data Visualization Workshop at IEEE VIS 2016, Baltimore, United States, October 2016.

[WHJ23] Wei Wei, Samuel Huron, and Yvonne Jansen. Towards autocomplete strategies for visualization construction. In *2023 IEEE Visualization and Visual Analytics (VIS)*, pages 141–145, 2023.

[WNT+20] Hsiang-Yun Wu, Benjamin Niedermann, Shigeo Takahashi, Maxwell J. Roberts, and Martin Nöllenburg. A survey on transit map layout – from design, machine, and human perspectives. *Computer Graphics Forum*, 39(3):619–646, 2020.

[Woo08] R. F. Woolson. Wilcoxon signed-rank test. In *Wiley Encyclopedia of Clinical Trials*, pages 1–3. 2008.

[WSK+19] Yun Wang, Adrien Segal, Roberta Klatzky, Daniel F. Keefe, Petra Isenberg, Jörn Hurtienne, Eva Hornecker, Tim Dwyer, and Stephen Barrass. An emotional response to the value of visualization. *IEEE Computer Graphics and Applications*, 39(5):8–17, 2019.

[YAD+18] Vahan Yoghourdjian, Daniel Archambault, Stephan Diehl, Tim Dwyer, Karsten Klein, Helen C. Purchase, and Hsiang-Yun Wu. Exploring the limits of complexity: A survey of empirical studies on graph visualisation. *Visual Informatics*, 2(4):264–282, 2018.

[Yuk15] Cem Yuksel. Sample elimination for generating poisson disk sample sets. *Computer Graphics Forum*, 34(2):25–32, 2015.

[ZSCM17] Bo Zhu, Mélina Skouras, Desai Chen, and Wojciech Matusik. Two-scale topology optimization with microstructures. *ACM Trans. Graph.*, 36(5), July 2017.

[ZVM08] Jack Zhao and Andrew Vande Moere. Embodiment in data sculpture: a model of the physical visualization of information. In *Proceedings of the 3rd International Conference on Digital Interactive Media in Entertainment and Arts*, DIMEA '08, page 343–350, New York, NY, USA, 2008.

Appendix

For all projects presented in this dissertation, supplemental materials were collected and published alongside the papers. This appendix gives a summary of the content of the supplemental materials, as well as a link to the online repositories.

HoloGraphs

The supplemental material for Chapter 3 contains the network diagrams for our *HoloGraph* of the Harry Potter dataset used in the case study, the codebase to create custom *HoloGraphs* from timesliced network data, 3D printable meshes to construct the slide holder, and a video demonstration. The online repository with the supplemental material can be found at `https://osf.io/4u2e9/files/`.

NODKANT

The supplemental material for Chapter 4 contains the questionnaires used in the study, an instruction booklet for the layout construction condition, an edge list for the free construction condition, used in the study, all data collected during the study, the codebase with scripts to create printable toolkits, instruction booklets, and a video demonstration. The online repository with the supplemental material can be found at `https://osf.io/tk3g5/`.

Squishicalization

The supplemental material for Chapter 5 contains the questionnaires for the initial study and the follow-up study, all data collected during both studies, the codebase for the tool to create custom *squishicalizations*, 3D printable files for the presented results in the paper, and a video demonstration. The online repository with the supplemental material can be found at `https://osf.io/35gnv/`.

Heart Machine

The supplemental material for Chapter 6 contains the questionnaires used in the study, the data collected during the study, and a video demonstration. The online repository with the

supplemental material can be found at https://osf.io/vxshu/.