Techniques for Physical Storytelling

by

Sarah Kushner

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy

Department of Computer Science
University of Toronto

Techniques for Physical Storytelling

Sarah Kushner
Doctor of Philosophy

Department of Computer Science
University of Toronto
2025

Abstract

Storytelling is what binds us together as humans. It's what entertains us, what moves our emotions, what preserves our cultures. Physical storytelling has a long history coming from traditional cave paintings and clay sculpture, and has developed into full experiences, such as theme parks and haunted houses.

This thesis explores how physical storytelling can be enriched through innovative techniques in floating sculptures, zoetropes, and animatronics. The primary objective was to enhance the audience's experience and broaden the accessibility of these storytelling mediums.

In the realm of floating sculptures, this research introduces an algorithmic approach to optimizing mechanical support structures, enabling the creation of visually immersive, walkthrough exhibits with concealed supports. This method was validated through the construction of physical examples, demonstrating its practicality in enhancing viewer engagement and interaction with floating art.

Augmenting zoetropes to be interactive, we add layers of narrative depth by integrating audio into the experience and taking advantage of the strobing light that creates illusion of motion. This research developed a new way of revealing hidden plot elements and engaging viewers, thereby expanding the storytelling capabilities of this traditional medium.

This thesis also addressed the potential of animatronics in educational contexts, particularly within K-12 settings. An affordable, versatile kit was designed to enable students to create and perform stories through papercraft puppetry and simple electronics. This approach aimed to foster cross-disciplinary skills and challenge the conventional boundaries between art and engineering, empowering students to become creators and storytellers.

The techniques developed in this thesis, combined with ongoing advancements in physical storytelling, pave the way for more complex, interactive, and accessible storytelling experiences. These innovations hold promise for more applications, from educational environments to immersive entertainment spaces, making storytelling a more inclusive and dynamic form of expression.

To my baby brother David, I miss you. I know you would be so proud of me.

Acknowledgements

Finishing this degree and writing this thesis would not have been possible without a number of people very important to me.

Thank you to my husband Tom, who has been extremely patient, supportive, and encouraging through incredibly hard times. You are the reason I still have faith in people. We have gone through so much together; whether it's challenging or exciting, there is no one I would rather do it with. Thank you to my older sister Jackie, who has been right by my side since she came into my life. I don't know what I would do without you. And your husband Eric is cool too. I could write another dissertation on exactly how much you mean to me. My cousin Lily is one of the strongest and smartest people I know. texting you everyday is something that keeps me going. I am so proud to call you my family.

Daniel, thank you for the high-level guidance through my first user study and for continuously asking "Does it answer your research question?" It means so much to me that you agreed to take me as your student and I am immensely grateful.

Paul, not only do I enjoy learning about 100 fun facts every time I talk to you, but I also value your engineering wisdom and patience in teaching me. Your crazy ideas and whimsy will continue to inspire me; you never seem to think any idea is impossible, and this way of thinking pushes my research to be better and strengthens my creativity in problem solving.

To my committee members Karan, and externals Tovi and Ann, thank you for taking time to read my thesis, give invaluable feedback, and help me along in my journey towards graduation. I hope to see you all at conferences in the future!

My time as a visiting student at UC Santa Barbara was life changing for me. Thank you to Nina Miolane and her postdocs Bongjin and David, students Christian, Fran, Mathilde, and Adele for believing in me and showing me that research can be fun. Thank you to my roommate and honorary little sister Ashley. Without you all, I probably would have quit my PhD.

Thank you to my undergraduate professors who made me love math: Sergio Zefelippo, Daryl Falco, and Dimitrios Papadopoulos; and who encouraged my interests in graphics and helped me get involved in research: Jarvis Thompson, David Breen, and Marcello Balduccini.

My sister Sarah White and other mom Cynthia have helped me in ways I can't express. Thank you for coming to Toronto to keep me on track and make sure I take breaks while writing this document. Thanks for being my neatest friend.

My Toronto friends have stuck by me through some of life's hardest moments, and have created some of life's most delightful moments: Sara, Alisa, John, Mengfei, Martina, Kieran, and Erin. My Philly friends back home have supported me from afar and made sure to see me when I visited: Sam, Jared, Dan, Julie, Jia, and Ruth. My online friends showed me that real bonds don't have to be in person: Maddy, Naoky, Adam, Jasiel, Don, Caroline, Dan, Macaroni, Stephen, Bates, Andrew, and Erik. My gymnastics coaches Zoia and Marlowe and therapist Cat have helped me grow mentally and physically, and have taught me to enjoy life more than they probably know.

I also want to mention my lab mates who I could always go to for advice: Michael, Nicole, Rahul, Oded, Christina, and Rinat; my lab mates in the NRG group: Bingjian, Bogdan, Jessi, and Dash; my lab mates who were in the thick of it with me: Mian, Vismay, Brenna, Silvia, Otman, and Derek; my undergrad and high school students: Niki, Evelyn, Arjun, Cindy, and Kevin; and my lab mates who motivate me with their hard work and cool research: Lawson, Gavin, Jacob, Selena, Sarah, Aravind, John, Ty, Jonathon, Parsa, Robin, Esther, Kinjal, Karran, Mengfei, Emilie, Zhecheng, Yuta, Chang Yue, Warren, Sue, Karthik, and Towaki.

John Hancock holds the Dynamic Graphics Project Lab together and it would not function without him. Thank you for making sure all the equipment needed for my experiments is available and working. Thank you for keeping the hardware room up to date and dropping everything to let me borrow a laptop when mine broke the day before my defense! Xuan Dam, thank you for dealing with my ridiculous reimbursement requests, for example, when I needed to buy expensive metal parts to build a giant zoetrope. Sara Lavoie, Joseph, Simon, and Nico, thank you for the IT support, endless admin work, and fun conversation in the hallways.

My parents have poured love and sacrificed time and money for me to get to where I am today. My dad taught me both how to draw and how to create webpages, cultivating my love for art and coding which eventually led me to computer graphics research. My mom taught me all the important life skills I know, and I would not be a functioning member of society without her.

Lastly, I would like to acknowledge the pets that have been through my PhD with me: to my rabbits Jack and Cinnamon Louise Bun and cat Anchovy, I miss you all so much. To my new baby, Dr. Sardine Pepper, I love you.

Contents

List of Tables

List of Figures

CHAPTER 1

Introduction

Once upon a time, a PhD student set out to identify gaps in the world of physical storytelling and offer creative solutions that change the way authors and audiences alike fundamentally interact with the stories they tell and experience. This thesis aims to push the boundaries of the art of physical storytelling by addressing areas of improvement in the way we currently tell stories. There are natural limits of bringing existing stories, be they text, image, film, or ideas still purely in the imagination of the storyteller, into physical space. For example, a storyteller may wish to fabricate a story and build a world which contains characters or objects that must appear as if they are floating in midair. In physically building a scene with floating objects, the artist would need a way of holding those objects up under the force of gravity. This is just one instance where physical storytelling faces limitations, and there will be more to discuss in the following chapters. We aim to invent technologies which mitigate these shortcomings in a creative way. In each physical medium we tackle, we challenge the dynamic of the relationship between the story and the people who make or consume it, allowing for more complex stories to be told. But first, let us define our terms.

1.1 What are physical stories and why do we tell them?

Storytelling is an important part of humanity. Humans use stories to convey information, preserve memories and legends, evoke emotion, spark imagination, and learn to feel empathy for each other and the world around us [53, Chapter 5]. From the cave paintings of the Stone Age, to the oral traditions in ancient Greece, to the first books made with the printing press, to the 3D IMAX films we watch in theatres today, stories are an important part of virtually every culture [126]. Humans have been storytelling since the beginning of our existence and show no signs of stopping. In fact, I would argue that with the advent of the internet, storytelling is more democratised and public than ever before. Anyone with an internet connection, freedom of expression, and creativity can rapidly disseminate their ideas to a potentially large audience [145].

Figure 1.1: Andrew Wyeth's 1948 painting *Christina's World* [144] tells a story where viewers are left to interpret many details. The audience can relate to the subject of the painting without knowing her exact backstory. The colours, the woman's pose, the composition of her looking toward the houses in the distance all contribute to the mood and potential meanings.

One might define a story as a narrative which contains all the elements an audience or elementary school English teacher would expect: a clear beginning, middle and end, developed characters, setting – a time and a place, a call to action, and a plot following a problem that resolves in the end. However, not every story explicitly contains every component, and this does not make it less of a story but rather a condensed message with room for audience interpretation. For example, a painting may convey a narrative, but some parts have to be implied or imagined by the audience (see Figure 1.1).

Many artists leave parts of the meaning of their works up to interpretation, and this too is part of the art [121]. An essential element of storytelling lies in the audience's role as an active participant in constructing meaning. This is not a shortcoming but a feature, as humans naturally fill in these gaps to create meaning. The Heider-Simmel animation study famously demonstrated this tendency, where simple geometric shapes moving on a screen were interpreted by viewers as characters in an emotional narrative [54]. This innate drive to assign meaning highlights that storytelling is not simply a process of transmitting information but one of co-creation. Discussions among audiences further amplify this process; for example, interpretations of an art piece evolve as viewers share insights, challenge assumptions, and build a collective understanding [17]. Stories also do not have to be long and detailed to be understood and have impact. Consider this short, sweet, and sad story:

> For sale: Baby shoes, Never worn.

It paints a picture of the imagined backstory of the character selling their brand new baby shoes [140]. Importantly, the story also stirs emotion in the audience.

There are many methods of telling a story. Broadly these methods can be categorised as oral, written, and visual. Many storytellers often employ a combination of methods to get their point across. Oral storytelling consists of telling a story with a voice and gesture; it has likely been around since the beginning of language. It can come in many forms such as speeches, poems, and songs, that can be paired with a dance or performance, like puppetry. Oral storytelling has helped humans pass down cultural history, teach lessons, and entertain for generations [113]. One example of oral storytelling would be a stand up comedy show: a comedian telling a story to an audience, which brings the audience joy and a sense of belonging [16]. Written storytelling uses written words to tell a story. This can include media such as newspaper articles, novels, poetry, ancient scrolls and much more. Visual storytelling uses images, symbols and other visual media; over time it has evolved from pictographs and icons on cave walls to real-time 3D video games. Visual media like illustration, photography, comics, and paintings are effective as storytelling methods because they catch the eyes of the target audience, holding their attention while both sharing its message and still allowing the viewers to inject their own point of view.

Over time, new storytelling mediums have developed, each uniquely suited to the society and technology of its era, but the essence of storytelling – a shared experience that speaks to its audience – remains unchanged. In the 20th century, media theorist Marshall McLuhan asserted that "the medium is the message," suggesting that the characteristics of a medium shape the audience's experience as much as, if not more than, the content itself [91]. This idea becomes central to understanding physical storytelling, where the medium's tangibility, interactivity, and spatiality uniquely contribute to the story it tells. The concept of story for our purposes is a flexible definition. The point of a story is that it takes the audience's mind on an adventure.

Physical storytelling requires additional nuance in its definition. A physical story, generally speaking, is a story told through the placement and sometimes movement of physical objects in an environment. They may or may not include audio or tactile aspects; the same principles of leaving room for viewer interpretation through its length, level of detail, and composition apply as they do for a non-physical story. What makes defining a physical story tricky is the fuzzy boundary between 2D and 3D objects, tangible objects and objects not meant to be touched, and the inclusion or exclusion of digital or computational elements. In any case, the point of a physical story is no different from a non-physical story.

1.2 Why physical storytelling?

The physical medium of a story is distinct from a non-physical story because it anchors the experience in a shared, tangible space. While an audiobook or movie is experienced digitally, physical storytelling unfolds in person, often involving three-dimensional elements and tactile interactions. A sculpture, for instance, offers physical depth and presence, and even an oil painting, though two-dimensional, exists as a tangible object in space. Some media, such as printed books or pop-up books, sit at the boundary of physical storytelling, inviting contemplation of whether they function as physical stories based on how their materiality affects the audience's experience. Through a physical lens, a story becomes an immersive experience, where audiences are not only spectators but participants in a shared space. This echoes McLuhan's idea, as the physical medium's form becomes integral to the story's impact, amplifying its message through sensory experience and shared presence.

The history of physical storytelling is vast, with examples ranging from the dioramas and shadow puppetry of early cultures [24] to modern theme park rides (e.g., Pirates of the Caribbean boat ride at Disneyland [62], light and sound shows (e.g., "Sleep No More", a promenade theatre walkthrough telling of Shakespeare's Macbeth [9]), and other immersive installations (e.g., Little Canada, a museum featuring crafted 1:87 miniature scale landmarks of Canada [14]). In contemporary settings, escape rooms, haunted houses (e.g., "Terror Behind the Walls", a spooky experience at an abandoned prison [109]), and interactive exhibits (e.g., Meow Wolf, an exploration of the multiverse through a

Figure 1.2: Karine Giboulo's Housewarming Exhibit at the Gardiner Museum, Toronto, Ontario, Canada [45]. A tiny gardener and her dog harvest vegetables in the drain of a sink.

grocery chain called Omega Mart [40]) continue this tradition, allowing stories to unfold as audiences engage with their physical surroundings. Each of these media underscores that the medium is not just a vessel for content but an active, shaping force that defines the story's resonance and meaning.

Guiding the audience's experience is fundamental to physical storytelling, ensuring key elements are seen and experienced as intended. Photographers and painters often lead the viewer's eye through composition, lighting, and movement. Video games and museums use methods of indirect control by limiting viewer choices, careful interface design, and establishing a goal [122, Chapter 16]. Additionally, more technical methods like subtle gaze direction have been explored. For instance, researchers have developed techniques using brief, peripheral image modulations to draw attention without the viewer consciously noticing, leveraging properties of human vision and phenomena like saccadic masking [8].

These principles can inform physical storytelling, where the deliberate arrangement of visual and interactive elements guides the audience through the narrative, shaping how they engage with the story and the space it occupies. This connection between the audience and the physical world creates a unique sense of presence and immediacy. There is something special to experiencing a story in person. In an increasingly digital age, consuming stories through a medium other than a screen can be refreshing and novel. A good example of a physical story is the exhibit in 1.2.

The Housewarming exhibit was all about artist Giboulo's experience during lockdown. Her observations about the state of the world, the anxiety around COVID-19, and her feelings of isolation due to her illness were captured in an entire house worth of sculptures. The audience could learn the artist's perspective, relate to the other people perusing the gallery around them, and simultaneously be in awe of the craftsmanship of the realistic miniatures within life-size household appliances.

In this thesis we focus on three specific media used to tell physical stories: floating sculptures – objects appearing to levitate, zoetropes – spinning animation devices which induce the illusion of motion of many still frames, and animatronics – scenes of mechanically moving and speaking robotic puppets. Each medium has been around for quite some time – since the 1930's [19], 1860's [85], and 1960's [61] respectively – although the principles and predecessors of these art forms have existed for much longer.

1.3 Overview

Chapter 2 describes a problem inherent to telling a story in a medium like a physical walkthrough experience or a stop-motion film containing floating actors or objects. How can existing mechanical engineering optimization techniques be adapted to fit our needs of not only satisfying force and torque balance constraints on the objects while also being hidden from view? Normally, an artist would have to manually place support structures in order to hold the object up. Unsightly supports distract from the story, so we provide a method of generating supports hidden from the viewer. We offer an automatic algorithm which invisibly supports the floating objects with an as-hidden-as-possible network of rods and wires from a given distribution of viewpoints. Can we validate our method via constructing examples of wire and rod structures in real life?

In Chapter 3, our goal is to expand the domain of stories possible to tell through a zoetrope. As stated, a zoetrope is an animation device consisting of a sequence of still images or statues placed around a cylinder. When a viewer sees the cylinder spinning at a particular rate, the still frames appear to move. We explore ways to increase the sophistication of the stories that would otherwise be limited to a silent, passively watched periodic animation. Because zoetropes spin, stories in this medium are constrained to be very short and repeat over and over. More specifically, in this thesis, we explore ways to make zoetropes interactive. How can light and audio be utilised to enrich a zoetrope story? How can we reveal hidden elements in a scene to engage the viewer and subvert their expectations and challenge their assumptions? We want to take advantage of the strobing light mechanism to add texture and depth to a zoetrope story.

In Chapter 4, we aim to give the ability to make animatronic stories to young children. Animatronics are robotic puppet characters that are programmed to act out a story. The art of making animatronics is not necessarily accessible for everyone since it requires mechanical construction of the puppets, a way to control their movements to look realistic, and synchronized audio to the motion. This medium may seem out of reach for non-experts, but we aim to democratise it and prove its usefulness in education. How can we provide K-12 teachers and students with tools to tell stories through animatronics? How can we engage students in storytelling to learn other subjects in school? As stated, a story can be a very broad term, which makes animatronics a versatile tool to use in the teaching of almost any subject. Finally, because animatronics is both an art *and* a science, we can use it to show kids that they are capable creators across disciplines. Many young kids are placed into boxes that label them either artists or engineers, which prevents them from thinking they can do the other. How can we use animatronics to encourage kids to change their perception of their own abilities?

1.4 Contributions

This thesis explains my contributions of the techniques and tools for telling stories in the physical world. Chapter 2 motivates the desire to make objects appear as though they are levitating in the context of telling a story through them, and provides a convenient way of building them. Chapter 3 tells the history of the animation device, the zoetrope, and shows that with our tools and methods, we can tell more complex, layered stories using audio and user interaction. Chapter 4 recounts the history of the art and science of animatronics, and demonstrates that our tools for making this activity accessible in the context of education succeed in enhancing learning outcomes and creative expression in a K-12 setting. The text and figures in this thesis have significant overlap with my publications over my PhD studies, including:

- 2021 Eurographics paper: Levitating Rigid Objects with Hidden Rods and Wires [80]
- 2022 UIST Demo: Interactive Zoetrope with a Strobing Flashlight [79]
- 2024 IEEE Frontiers in Education conference paper: Papertronic Puppets: Teaching STEM and Storytelling Through Creative Construction [81]

Although the technical innovations in this thesis are emphasized, conceptually my work pushes the boundaries and considers how these innovations will have a lasting impact on the field in terms of immersion, depth, and accessibility of physical storytelling as a whole. Now, let's explore the enhancement of physical storytelling through clever application of mechanics and electronics!

CHAPTER 2

Floating Sculptures

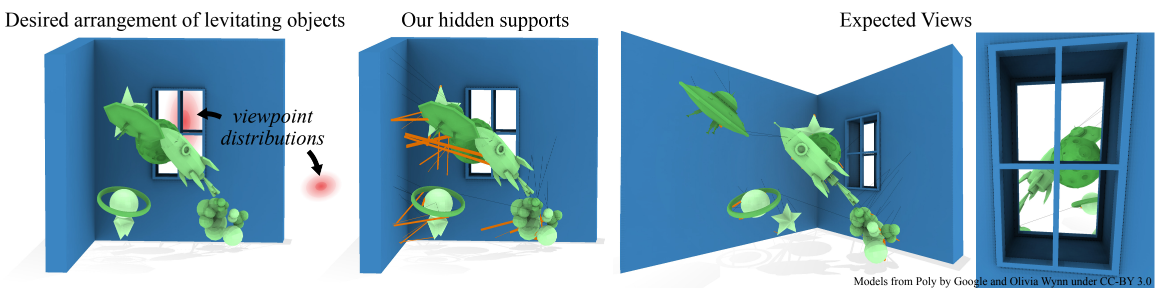

Figure 2.1: Our optimization finds hidden supports to hold rigid objects (green) in their locations despite gravity. Rods (orange) resist tension, compression and bending, while wires (black) resist tension. Supports connect between objects or to the input support surface (blue). Rods are *hidden* behind occlusions in the scene for a possibly disconnected distribution of viewpoints (red) provided by the user. Here, a collection of space-themed objects seemingly hover in the corner of a room. The supporting truss is hidden from the front *and* through the window.

2.1 Introduction: Creating Invisible Support Structures in Physical Walkthroughs

Levitating objects are visually compelling and commonly found in artistic sculptures, film and theatre set design, promotional displays, and museum exhibits (see Figure 2.2 and Figure 2.4). This effect is especially impressive if the support structure can be hidden from the observer, removing its unsightly distraction and perhaps even giving the impression that the objects in the arrangement are magically floating in space (see Figure 2.1). Achieving this is a non-trivial task. Physical stability requires a balance of force and torque for each rigid component of the scene. This is readily achieved using many strong, thick struts, but their geometry and scene placement is likely to compete for visual attention with scene objects, or worse, visually obscure objects in the scene (see Figure 2.3). Hiding these supports by removing or thinning too many struts, on the other hand, will sacrifice physical stability. Thin wires can sometimes be used to hang objects, but wires only resist tension so they alone can not handle situations that are not supported purely from above.

In this chapter, we propose modeling the problem of hidden support structure generation for *levitating objects* as a form of topology optimization. We present a novel convex optimization based on the well-established *ground structure method* from architecture and engineering. The input to our method is an arrangement of objects in their desired locations and orientations and the distribution of views from which the scene will likely be observed. Our output is a collection of *rods* and *wires*, described by their required thicknesses and attachment points on the input rigid objects, and the supporting structural element (e.g., wall or ceiling). Our rods model tension, compression and bending resistant materials (e.g., wooden dowel rods or steel beams). Our wires model tension only (e.g., fishing line or steel cables).

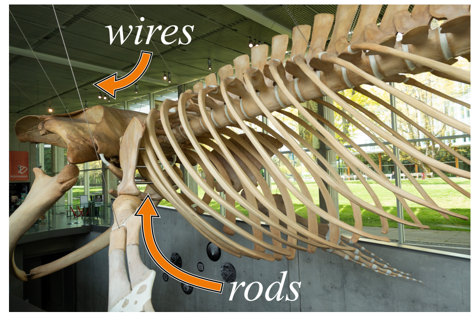

Figure 2.2: A blue whale skeleton floats with support of wires and internal rods.

Unlike computer graphics or virtual reality (VR) where physical laws can be bent or broken, support structures in real scenes are only meaningful if physically valid. Therefore, we enforce physical validity in our optimization as a hard constraint: namely that the rigid objects should achieve force and torque equilibrium and that stresses on rods and wires do not exceed material-dependent yield limits. For ease of assembly, cost of manufacturing, and visibility considerations, we prefer support structures composed of a small number of thin, less visible supports. We model these criteria with a sparsity-inducing cost function defined as a sum over a densely connected graph

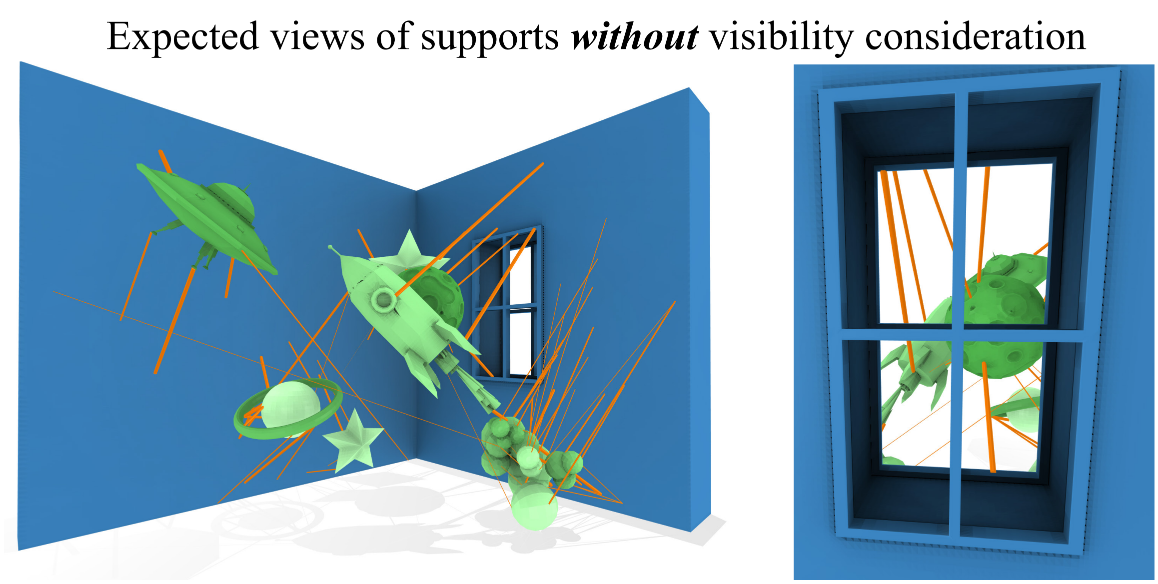

Figure 2.3: Without our visibility term, optimal rods may be an unsightly distraction.

Figure 2.4: Levitating objects have inspired such artworks as a sculpture of border guard Conrad Schumann jumping (left), an enormous stage display of playing cards by Es Devlin for the Bregenz Festival (middle), and Chiharu Shiota's installation where white dresses float overhead (right).

of edges (i.e., the *ground structure*).

Treating the cross-sectional area of each edge as the primary optimization variable, the traditional ground structure method optimizes the total volume (linear in the areas since lengths are predetermined) and enforces force balance at point loads, by measuring linearized axial tension and compression forces from each rod, subject to yield limits, expressed as linear inequalities in the unknown cross-section areas and axial stresses of the rods. The result is a linear program whose solution — like many L_1 or Lasso problems — is sparse (most areas are exactly zero), and often agrees exactly with the NP-hard selection problem (picking the smallest valid subset of edges).

We augment the traditional ground structure method to support embedded rigid objects (via linear static equilibrium equations) and account for bending resistance of rods (via a simple linear shearing model derived from proportionality assumptions). We introduce a visibility objective function that is also linear in the unknown edge areas and relies on efficient Monte-Carlo based precomputation. Thus, the optimization remains a (convex) linear program and solutions can be extracted efficiently (in usually less than a minute).

Our experiments satisfyingly confirm that under many conditions structurally valid supports are lurking just out of sight: the space of physically valid supports is vast and finding a completely occluded arrangement is often possible. We demonstrate the effectiveness of our method across a wide variety of test scenes and prototypical use cases.

2.2 Related Work: Mechanical Optimization & Structural Design

Our work sits within the larger literature of computational fabrication, construction and assembly. These subfields are rich and vast, so we focus on previous works most similar in methodology or application.

Previous algorithms exist to make objects stand [112, 134], spin [7] or hang from wires [87]. These works modify the input objects by redistributing mass or changing their shape to achieve the desired goal. In this chapter, we explore a complementary contract with the user — how to anchor objects in the environment *without* changing the objects themselves. We do not assume that objects were fabricated in a particular manner (e.g., 3D printing).

Our approach may be categorized with other structural optimizations for a prescribed static load scenario (i.e., ignoring inertial forces). Recent works increase the stability of fragile objects by adding new structural elements [151, 129, 28]. For example, Stava et al. add struts to 3D printed objects one-by-one as part of a large optimization loop and use a volumetric simulation as validation. Their strut selection includes an ambient occlusion visibility term, but they do not consider the problem of selecting an optimal set of supports for rigid objects under prescribed viewing conditions. Other methods have considered the interactive design of rod structures [110, 77, 22, 66] with varying degrees of physical feasibility checking or optimization in the system.

We model the problem of hiding support structures as a form of topology optimization [86]. The general idea of topology optimization is to prune away material from the volume around the input objects or load conditions. The resulting geometries typically have interesting topologies/connectivities that would have been difficult to determine *a priori*. Methods that determine the material occupancy of each voxel in a dense grid are well suited for 3D printing and milling (e.g., [143]), but will in general produce geometries composed curved and varying thickness elements. Our method instead belongs to the class of *ground structure* methods [35], which output a discrete collection of (straight) elements from an initial over-connected graph of candidates (see Figure 2.7). Methodologically we follow most closely the stress-based formulation of Zegard et al. [148], and utilize the thesis of Freund [42] as a reference. Ground structure-like methods have been applied for designing everything from buildings [149] and glass shell structures [41] to construction supports [32] to 3D printable models [138, 69, 60] to cable-driven automata [92]. The standard ground structure method considers only axial forces. These methods have been applied for rigid structural elements and adapted to special cases like tensegrities [111, 27].

We use a ground structure approach to model the novel problem of creating hidden structural supports from complex viewpoint distributions. Crucially, our method supports structural elements that resist compression, tension and bending forces, as well as wires, without resorting to the non-linear constitutive models or volumetric meshing of prior work [129, 110, 60]. Our method trivially couples the structure to the rigid objects it supports, correctly accounting for both linear forces and torques, without resorting to displacement-based mechanical formulations (e.g., [111]). This allows us to formulate our problem as a linear program which can be solved efficiently.

We draw inspiration from algorithms for appearance-driven optimization. For instance, Schuller et al. introduce the problem of generating appearance mimicking surfaces from a specified viewpoint [123]. Several works seek to create 3D shapes that take the form of a set of 2D shapes from corresponding viewpoints or cast the 2D image under certain lighting conditions [97, 58, 124]. Others use viewpoints to create optimal perceptual experiences, for example in 3D printing support structures [150] or in skyscraper design [34].

2.3 Method: Developing an Algorithm for Hidden Support Generation

The input to our method is a scene comprised of K rigid objects oriented and positioned in space, a fixed support surface (e.g., wall or ceiling), and a distribution of viewpoints (e.g., discrete set of

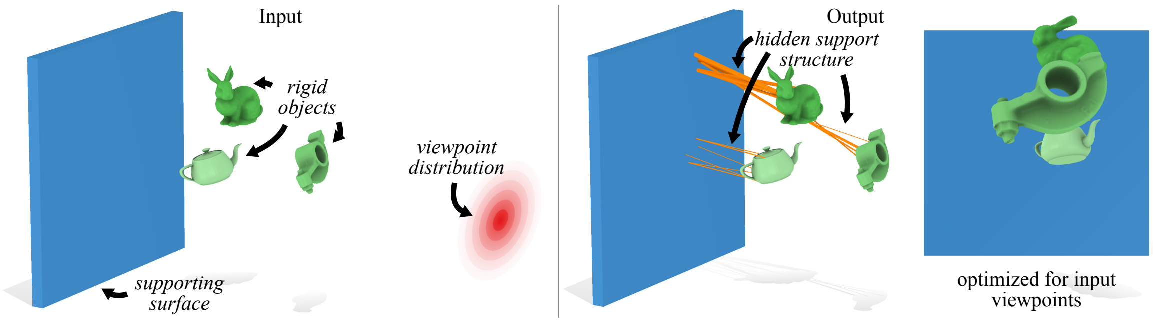

Figure 2.5: The input to our method is a scene composed of many levitating rigid objects. The output of our method is a collection of rods tucked away behind object occlusions, holding each object in force and torque equilibrium under gravity.

positions or sample-able probability density function defined on a surface) The output of our method is a *supporting structure* composed of a small set of *rods* and *wires* connecting rigid objects to each other or the supporting surface. Our method ensures that this structure holds the input objects in their prescribed positions and orientations, counter-balancing the force these objects experience due to gravity. Our method optimizes the size and placement of the structure to minimize its overall volume and its visibility with respect to the input viewpoint distribution (see Figure 2.5). Before describing our optimization, we define our physical model and how we measure visibility.

2.3.1 Rigid Body Equilibrium

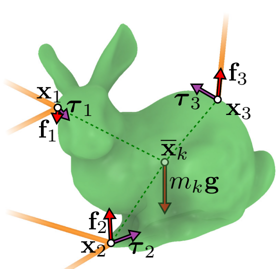

The rigid objects in our scenes experience forces from gravity and at the points of attachment to the supporting structure. To hold a rigid body at rest, we must maintain force and torque equilibrium:

$$\sum_{i \in V_k} \mathbf{f}_i = m_k \mathbf{g}, \tag{2.1}$$

$$\sum_{i \in V_k} \underbrace{(\mathbf{x}_i - \overline{\mathbf{x}}_k) \times \mathbf{f}_i}_{\boldsymbol{\tau}_i} = \mathbf{0}, \tag{2.2}$$

where m_k, $\overline{\mathbf{x}}_k$, and V_k are the mass, center of mass, and set of attachment points of the kth object, respectively, and $\mathbf{x}_i, \mathbf{f}_i, \boldsymbol{\tau}_i$ are the 3D position of the ith attachment point and corresponding force and torque vectors, respectively.

2.3.2 Rods

We assume our support structure undergoes negligible displacement, affording a linearization of the internal forces at play. For stiff rods, we follow the linearized tension and compression model of [148, 42], which introduces a *signed* scalar value per rod $c_{ij} \in \mathbb{R}$ with units Newtons describing the force in the axial direction parallel to the rod. Assigning an arbitrary direction to the rod ij between endpoint positions $\mathbf{x}_i$ and $\mathbf{x}_j$, then the axial force contribution at endpoints i and j are the product of this scalar c_{ij} by the rod's tangent unit direction $\hat{\mathbf{t}}_{ij} = (\mathbf{x}_i - \mathbf{x}_j)/\|\mathbf{x}_i - \mathbf{x}_j\|$:

$$\mathbf{f}_i \mathrel{+}= c_{ij}\hat{\mathbf{t}}_{ij} \quad \text{and} \quad \mathbf{f}_j \mathrel{-}= c_{ij}\hat{\mathbf{t}}_{ij}. \tag{2.3}$$

Previous methods (e.g., [148, 42]) rely solely on tension and compression and ignoring the rods' resistance to *bending*. This is a reasonable assumption in architecture where loads are large relative to the rod's bending strength. Ignoring bending requires that the rods are thicker and thus more visible (see Figure 2.10). This is at odds with the intuition that light loads can be held up with a single bending-resistant rod. In reality, a single rod with finite thickness can apply a *distribution* of forces over its non-zero area contact surface. Since the force is applied at more than one point, torque balance is also possible. Unfortunately, a volumetric rod model couples the unknown rod diameters and forces non-linearly.

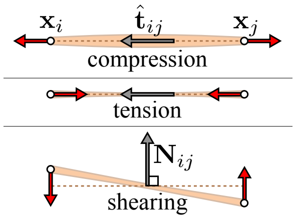

To maintain the linearity of our system but also account for bending, we introduce a linearized shearing model to account for resistance in the normal direction (see Figure 2.9). For each rod ij, we introduce an arbitrary orthonormal basis $\mathbf{N}_{ij} \in \mathbb{R}^{3\times 2}$ for the 2D space orthogonal to the axial direction. We introduce a two dimensional parameter $\mathbf{q}_{ij} \in \mathbb{R}^2$ with units Newtons describing the force on the rod in the two normal basis directions. Shear force contributions are equal and opposite at either end of each rod:

$$\mathbf{f}_i \mathrel{+}= \mathbf{N}_{ij}\mathbf{q}_{ij} \quad \text{and} \quad \mathbf{f}_j \mathrel{-}= \mathbf{N}_{ij}\mathbf{q}_{ij}. \tag{2.4}$$

Following previous methods [148, 42], we model failure catastrophically. If the stress due to tension, compression or bending exceeds a material-dependent fixed threshold we declare that the rod has exploded (or at least moved too much) and is no longer feasible. These yield stresses can be prescribed for each rod ij and can be related directly to the non-negative rod cross-sectional area $a_{ij} \in \mathbb{R}_{\geq 0}$ and the force parameters introduced above. Namely, we require the following *convex* inequalities to hold:

$$-\sigma^t_{ij}a_{ij} \leq c_{ij} \leq \sigma^c_{ij}a_{ij} \quad \text{and} \quad \|\mathbf{q}_{ij}\| \leq \sigma^s_{ij}a_{ij}, \tag{2.5}$$

where $\sigma^t_{ij}, \sigma^c_{ij}, \sigma^s_{ij}$ are the tension, compression, and shearing stress thresholds, respectively. For common rod materials, we find that $\sigma^t_{ij} \approx \sigma^c_{ij} >> \sigma^s_{ij}$. Although σ^t and σ^c values for specific materials (e.g., pine wood) can be found in reference books, in our experience all of these parameters should be empirically estimated, especially when working with low-end materials from the hardware store.

2.3.3 Wires

A special case of our model is a *wire*, which can be thought of as a tension-only rod. A wire ij has zero resistance to bending and compression (i.e., $\sigma^c_{ij} = \sigma^s_{ij} = 0$) and very high resistance to tension (i.e., $\sigma^t_{ij} >> 0$). Wires made of strong material such as braided steel can be very thin (near invisible) while maintaining high strength. Our method will allow a mixture of tension-compression-bending rods (e.g., wooden dowels) and

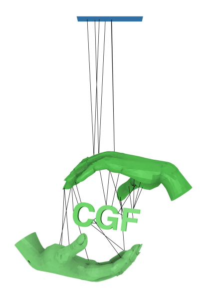

Figure 2.6: Wires only.

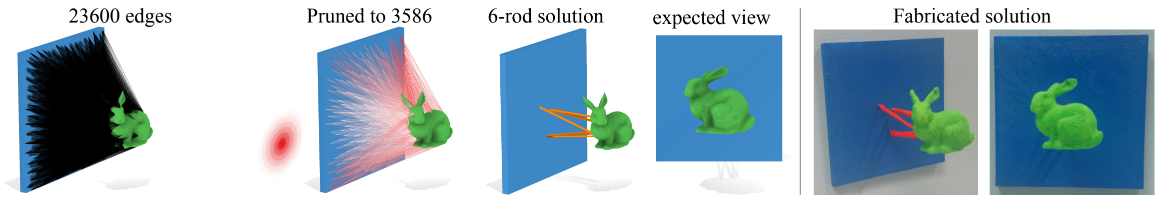

Figure 2.7: Our method constructs a over-connected ground structure of candidate edges (left) then immediately prunes edges that intersect the scene (middle) and finally extracts a small number of hidden rods. Savings from pruning can produce 10× performance improvements.

tension-only wires (steel wires), see Figure 2.11. As special case, we can limit our optimization to consider only wires, resulting in a hanging optimization (see Figures 2.6,2.14). Wire-only solutions require support from above the arrangement's center of mass.

2.3.4 Visibility

We define the *expected visibility* of a rod as function of the input viewpoint distribution, occlusions due to the scene, the rod's position and orientation and its unknown cross-sectional area. For a rod ij, its expected visibility v_{ij} is:

$$v_{ij} = \int_{\mathcal{E}} p(\mathbf{e}) \int_{C_{ij}} r(\mathbf{e}, \mathbf{x}) \, d\Omega \, d\mathbf{e}, \tag{2.6}$$

where

$$r(\mathbf{e}, \mathbf{x}) = \begin{cases} 0 & \text{if the segment } \mathbf{ex} \text{ intersects the scene,} \\ 1 & \text{otherwise,} \end{cases}$$

where $\mathcal{E}$ defines the set of viewpoints and $p(\mathbf{e})$ is the probability density associated with the point $\mathbf{e} \in \mathcal{E}$, and C_{ij} is the surface of the cylindrical rod with cross-sectional area a_{ij} connecting endpoints $\mathbf{x}_i$ and $\mathbf{x}_j$, and $d\Omega$ is the differential solid angle at the corresponding integration point $\mathbf{x}$ subtended at the viewpoint $\mathbf{e}$.

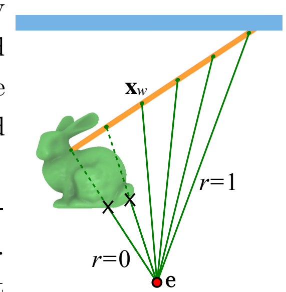

Measuring visibility according to solid angle correctly matches the intuition that the same size rod farther away from an observer is less visible. The outer integral is immediately recognizable as a soft-shadow or area-light source evaluation common in rendering. We can approximate this well by Monte-Carlo importance sampling over the viewpoint distribution. An analytic expression for the inner integral becomes unwieldy, so we instead opt for a simple approximation based on uniform quadrature, accounting for the orientation of the rod resulting in foreshortened projection (see Figure 2.3.4).

Rods are thin relative to the scene and spread of the viewpoint distributions, therefore we assume visibility to be constant in the normal directions of the rod. Our discrete approximation of the expected visibility is thus a double sum over n_u points sampled according to the input probability

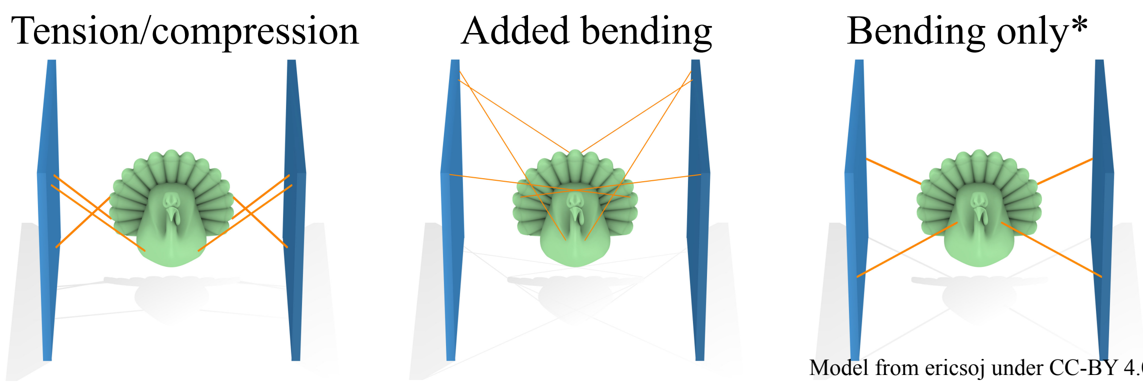

Figure 2.9: An enormous turkey levitates between two buildings using tension and compression resistant rods (left). Adding bending resistance affords a less voluminous solution (middle). Restricting the ground structure to only include edges perfectly intersecting the center of mass (*) admits a bending only solution (right).

density function and n_{ij} points sampled along the rod:

$$v_{ij} \approx \sqrt{a_{ij}}\, \underbrace{\frac{1}{n_u\sqrt{2\pi}} \sum_{u=1}^{n_u} \cos^{-1}\left(\frac{(\mathbf{x}_i - \mathbf{e}_u)\cdot(\mathbf{x}_j - \mathbf{e}_u)}{\|\mathbf{x}_i - \mathbf{e}_u\|\|\mathbf{x}_j - \mathbf{e}_u\|}\right) \sum_{w=1}^{n_{ij}} r(\mathbf{e}_u, \mathbf{x}_w)}_{g_{ij}},$$

where we collect the terms that do not depend on a_{ij} into a single non-negative scalar per-rod, $g_{ij} \in \mathbb{R}_{\geq 0}$. In this way, the *squared visibility* of each rod becomes a linear function of the cross-sectional area: $a_{ij} g_{ij}^2$. Because *wires* are so thin compared to *rods*, we happily set $g_{ij} = 0$ for *wires* and avoid their visibility precomputation.

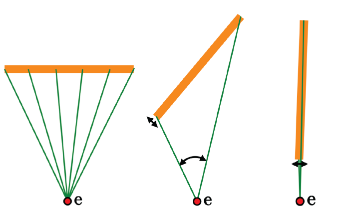

Figure 2.8: We compute the *projected* visible area of the rod which is a function of the angle between the vectors from the viewpoint to either endpoint.

To generate the n_{ij} samples on edge ij, we subdivide the edge until all segments are less than a given scene-dependent length threshold (e.g., 0.1 meters for the bedroom scene in Figure 2.1) and then use the segment barycenters as samples (typically 10-100 samples per edge). Segment queries can be computed in parallel.

2.3.5 Ground Structure

The space of physically feasible supporting structures is high-dimensional and a mixture of discrete variables (e.g., how many rods? connecting between which objects?) and continuous variables (e.g., where rods attach to each object?what are the rod thicknesses?). Navigating this space to find a *globally optimal* solution is difficult. In response, the ground structure method (e.g., [35, 107, 148, 42] makes the problem tractable by rephrasing the problem into *selecting* a discrete subset of support elements from an intentionally dense yet finite set of candidate elements. This candidate set is referred to as the "ground structure."

In our case, we generate a ground structure of candidate rod and wire elements by Poisson disk sampling [147] all rigid objects and the support surface and then connecting all possible pairs of

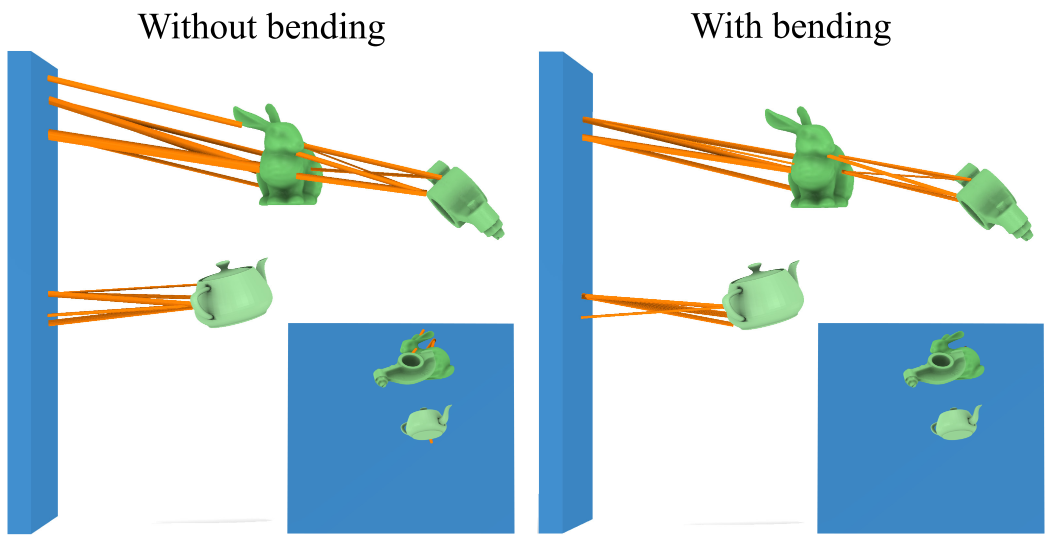

Figure 2.10: The addition of our linearized bending term yields sparser, less visible support structures by more accurately modelling the strength of the rods. Rods can also connect between objects rather than just to the support surface.

points from different sources (e.g, for a single rigid object this forms a bipartite graph with the supporting surface, see Figure 2.7). For each edge in this graph, we label it as a "rod" or "wire" (and possibly create duplicate copies so edges appear as both types). We can discard a *bad* edge ij if its attachment angle is self-penetrating or too obtuse (by checking if the rod vector dotted with the surface normal is below a threshold; $\hat{\mathbf{t}}_{ij} \cdot \hat{\mathbf{n}}_i < \cos\theta_{\max}$), if it intersects objects in the scene (by ray casting), or if its computed visibility coefficient is exceptionally high ($g_{ij} > g_{\max}$). We refer to the result as the *pruned* ground structure $\mathcal{G}$.

2.3.6 Sparse Optimization

The beauty of the ground structure method is that once the candidate set has been chosen, selecting the globally optimal subset can be phrased as an efficient convex optimization, in particular a linear

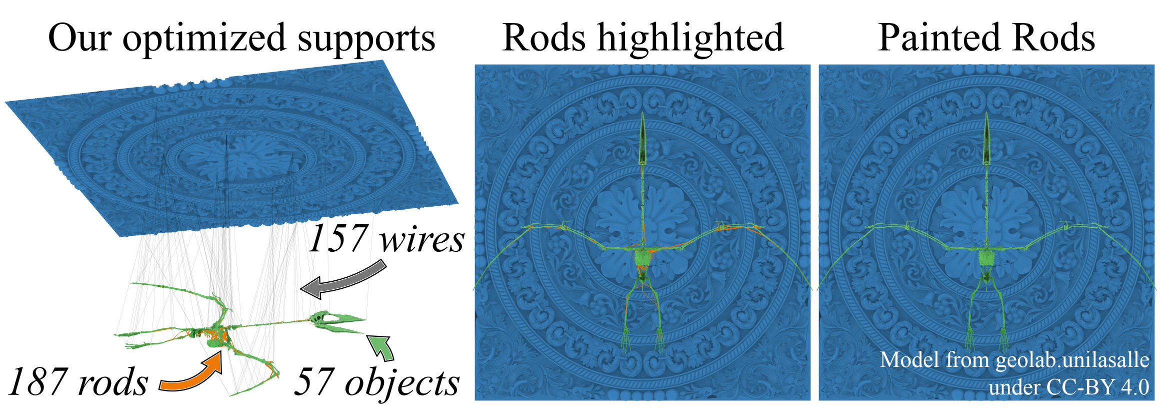

Figure 2.11: We show the rods in orange to demonstrate how hidden they are. But we show that the rods connecting the smaller parts (middle) can be painted to blend in with the ceiling (right).

program. In the classic method, the cost function to be minimized is the total volume of material spent on the support structure. Since all edge lengths are known once the candidate set is selected, this cost is a linear function of the yet unknown edge cross-sectional areas. It is important that this cost function is the *unsquared* volume, which can be thought of as the L1-norm of the vector of edge areas (weighted by edge-lengths), as opposed to the sum of squared per-edge volumes, analogous to the L2-norm. The L1-norm is *sparsity inducing* and under mild conditions will agree with the optimal solution of the selection problem, analogous to the L0-pseudonorm [20, 38]. As a result, the vast majority of edges in the solution will have exactly zero area.

In our case, we augment the total volume cost function with a least-squares visibility term to penalize choosing highly visible rods. Because our per-edge visibility measurement in Eq. 2.7 is linear in the square-root of the rod areas, this least-squares energy becomes linear in the areas.

The areas of the rods and wires are the primary unknowns. We introduce auxiliary variables c_{ij} and $\mathbf{q}_{ij}$ as described in Section 2.3.2 to facilitate writing our force and torque balance constraints (see Section 2.3.1). These variables are then coupled to the areas via the yield stress inequalities (see Eq. 2.5).

The resulting optimization is a linear program over the pruned ground structure $\mathcal{G}$ containing m candidate edges:

$$\min_{\mathbf{a},\mathbf{c},\mathbf{q}} \sum_{ij\in\mathcal{G}} a_{ij}(\ell_{ij} + \lambda g_{ij}^2) \tag{2.7}$$

$$\text{s.t.} \sum_{ij|j\in V_k} c_{ij}\hat{\mathbf{t}}_{ij} + \mathbf{N}_{ij}\mathbf{q}_{ij} = m_k\mathbf{g},\ \forall\, k = 1,\dots,K \tag{2.8}$$

$$\sum_{ij|j\in V_k} (c_{ij}\hat{\mathbf{t}}_{ij} + \mathbf{N}_{ij}\mathbf{q}_{ij}) \times (\mathbf{x}_j - \overline{\mathbf{x}}_k) = 0,\ \forall\, k = 1,\dots,K$$

$$-\sigma_{ij}^c a_{ij} \le c_{ij} \le \sigma_{ij}^t a_{ij},\ \forall\, ij \in \mathcal{G} \tag{2.9}$$

$$-\sigma_{ij}^s a_{ij} \le \mathbf{q}_{ij} \le \sigma_{ij}^s a_{ij},\ \forall\, ij \in \mathcal{G} \tag{2.10}$$

$$a_{ij} \ge 0,\ \forall\, ij \in \mathcal{G} \tag{2.11}$$

where we stack all a_{ij}, c_{ij}, and $\mathbf{q}_{ij}$ variables into vectors $\mathbf{a} \in \mathbb{R}^m$, $\mathbf{c} \in \mathbb{R}^m$, and $\mathbf{q} \in \mathbb{R}^{2m}$, respectively, and we introduce the user-controllable weighting term λ to balance between preference for volume and visibility minimization. For all examples shown, we use $\lambda = 10{,}000$.

We opt to replace the second-order cone constraint for linearized bending yields in Eq. 2.5 with the simpler coordinate-wise linear inequality in Eq. 2.10. This can be thought of as a conservative L_∞ approximation, and albeit coordinate system dependent, does not affect results and admits a faster linear program than a conic program in our experience.

The linear coefficients in the force/torque balance equations and linear inequalities (Eqs. 2.8-2.10) can be collected in large *sparse* matrices (see App. 2.6). Many efficient solvers exist for such large sparse linear programs; we use Mosek [3].

A solution is a guaranteed to exist as long as force and torque balance can be achieved. This could fail to happen for very sparse ground structures (e.g., less than six edges per object) or degenerate situations (e.g., all edges are parallel). Making our ground structure very dense (hundreds of thousands of edges) ensures that we never fail to find a feasible solution.

Scene	K	FULL	m	TIME	R	W
Parade Float	1	20K	10K	0.14	0	12
"Koons" Display	1	154K	22K	4.90	2	4
Ghost With Tail	1	11K	1K	5.01	5	0
Bunny/Teapot/Rocker	3	150K	21K	2.04	16	0
Bedroom	13	490K	46K	35.87	41	35
Zoetrope (1 frame)	1	469K	51K	33.77	4	0
Pterosaurus	57	10M	785K	608.01	187	157

Table 2.1: Timings in seconds (TIME) and numbers of rods (R) and wires (W) for each result with K objects. FULL and m are edges in the original and pruned ground structures, respectively. The "Pterosaurus" and "Parade Float" examples do not include the time for computing visibility, as it was not used in the LP objective.

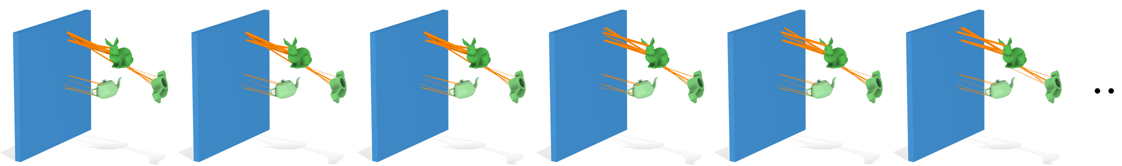

Figure 2.12: Our results depend on a randomly generated ground structure. Changing the random seed affects the precise result, but not qualitatively.

2.4 Experiments & Results

We implemented our algorithm in MATLAB using GPTOOLBOX [65] for geometry processing and MOSEK [3] to solve the linear program formulated in Section 3.6. Pre-computation of the integrated visibility, in our input scene is accelerated using the EMBREE [136] ray-tracer as interfaced by LIBIGL [67]. We report statistics and timings for the results in our paper in Table 2.1. All times are reported on a MacBook Pro with 3.5 GHz Intel Core i7 and 16GB of RAM. Visibility pre-computation is computed in parallel, but is still typically the bottleneck ($\approx 80\%$).

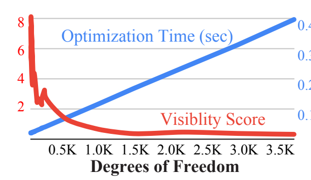

The number of degrees of freedom in the system is the size of the ground structure which generally scales quadratically in the number of objects $m = O(K^2)$, typically generated by taking all inter-object pairs over 10-100 Poisson disk sample points on each object. The inset graph shows the effect of increasing the degrees of freedom on Figure 2.5. While optimization time increases linearly with degrees of freedom, improvement to the visibility score of the solution reaches a point of diminishing return.

Starting with a dense ground structure leads to better qualitative results, but the exact positions of the samples do not drastically effect the hidden-ness of the result. Figure 2.12 shows how little the solution changes as a function of the ground structure sampling.

Pruning often significantly reduces the ground structure size and consequently, the number of degrees of freedom (see, e.g., Figure 2.7). Perhaps unsurprisingly, we typically experience a speedup the same ratio of original ground structure edges to pruned ground structure edges. The number of constraints in our optimization is six times the number of objects K. After pruning, MOSEK finds a solution for the above problem configuration within a few minutes.

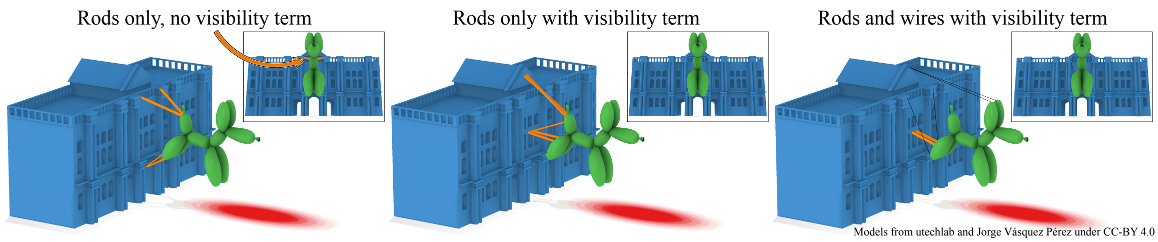

Figure 2.13: Applying the ground structure method to this example of a giant balloon hanging outside of a museum gives sufficient rods to support it, but they are visible. Using our visibility term in the optimization yields a support structure with rods hidden to the viewpoints. Allowing wires for tension and rods for compression, the result is a few thick but invisible rods and thin wires which hold the balloon in place.

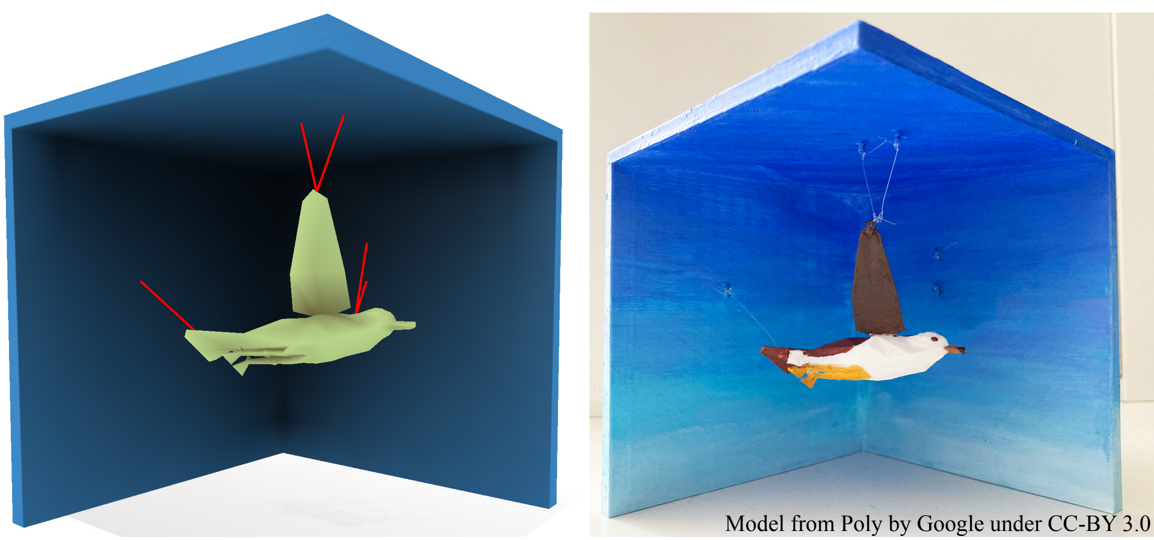

Figure 2.14: For a wire-only solution, we can save time by forgoing the visibility computation. Using fishing wire, we support the seagull in mid-air invisibly.

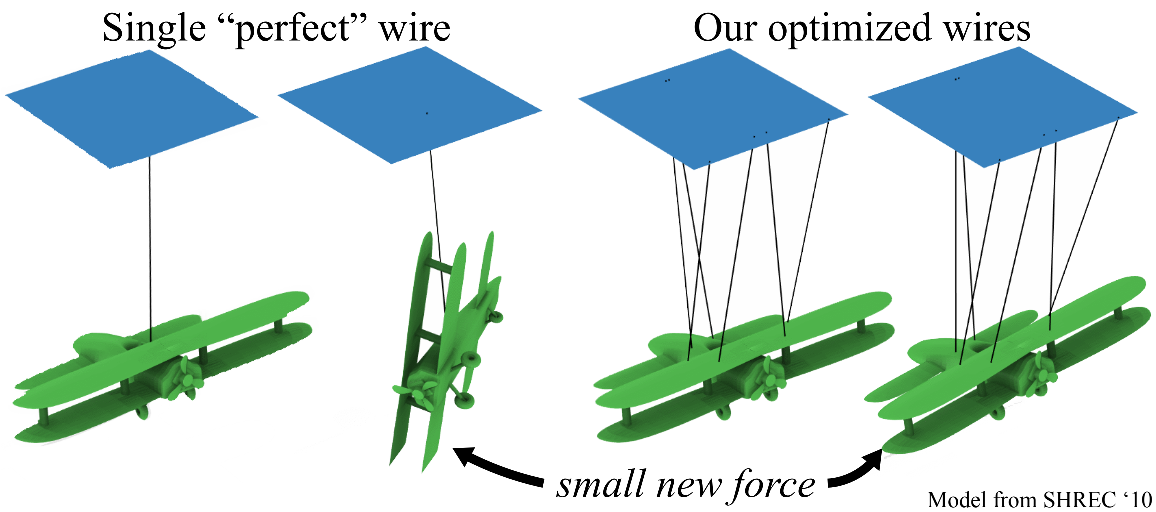

Figure 2.15: Applying a small new force to the plane held by a single wire causes undesired behaviour since a single wire attachment is not enough to balance the torque. Our method gives a 6-wire solution, exactly the number needed to balance force and torque.

Levitating 3D objects has a wide range of applications including scientific visualization, film and theater set design, home decor, anamorphic 3D art installations, as well as objects for zoetropes and 3D stop-motion animation. Each application has specific design requirements, and our algorithm is designed to enable the exploration of a number of aesthetic and structural parameters and design choices, which significantly impact the resulting solution. We elaborate on some of these design use cases.

Both scientific exhibits (see, e.g., Figures 2.1,2.2) and illusory art installations (see Figure 2.16) require an unobscured view of the levitating objects. While it is feasible for designers to hand-craft support structures from a single fixed viewpoint, the interplay between visibility and structural stability is quite complex for mutli-view distributions. In Figure 2.1, we show the ability of our algorithm to adapt its optimal solution to multiple viewpoint distributions.

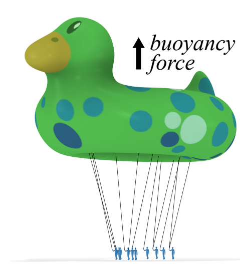

Our algorithm is able to holistically optimize the support structure using a mix of rods and wires. We color our rods bright orange for evaluation in this chapter, but in practice they can be further camouflaged by matching their appearance to the background or scene objects (see Figures 2.11,2.17). The choice of using a rod or wire is both aesthetic (as determined by a user) and functional. For example, supporting a levitating object with a wire would require a potential attachment points on the fixed surface or other levitating objects, to be vertically higher than the given object (see, e.g.,Figure 2.6). The inset figure shows a parade float suspended by optimized wires (the net force pointing upward due to buoyancy). Previous methods have considered hanging objects [112] or more generally mobiles [87] by placing a single support "perfectly" placed in alignment above the center of mass. While this strategy requires the fewest supports, it is an unstable solution (see Figure 2.15). Our method relies on random sampling of points in general position, typically producing multiple wires per hanging object, but resulting in a more stable configuration. Thus, in practice, we've found both our 3D

Figure 2.16: The model used for support attachments does not have to be the same one used for visibility. The ghost's head (green) has attached supports, while its tail (yellow) hides them.

printed and assembled results to be quite resilient to outside forces (including those from falls, heavy winds, and moving from one place to another).

Mounting objects off the side of a support such as a wall is best achieved with a mixture of wires and rods. Figure 2.13 shows a giant promotional display suspended in front of a contemporary art museum. We provide a symmetric dense ground structure and our optimization naturally finds a symmetric sparse solution.

The pterosaurus in Figure 2.11 has 57 separate bones and requires a complex support structure, acting as a stress test on our optimization. In practice, skeleton displays often pre-plaster-fuse bones to reduce the number of pieces (see spine of whale in Figure 2.2).

The idea of stop-motion animation and 3D zoetropes is over a century old [89], with modern examples including "Feral Fount" by Gregory Barsamian at the Museum of the Moving Image in Queens and the *Toy Story* zoetrope featured in Pixar's Museum Exhibit. The portrayal of levitating objects in this medium is particularly challenging. We demonstrate a prototypical result of a backflipping boy in Figure 2.19 by hiding supporting rods out of sight. For this example to be structurally stable both at rest and while spinning, we first find the optimal set of rods for each frame under gravity and then re-run the linear program on just these rods subject to centripetal forces. The final rod thickness are the maximum over the two solves. Incorporating more elaborate multi-load handling (cf. [42]) is left as future work.

2.5 Future Work: Support Generation Design Tools

Our rod model includes linearized tension, compression, and bending forces. Like many past methods, we do not handle the self-weight of the rods by assuming that the force of gravity is much larger than the force of the rods on themselves. This is a trivial addition of gravity forces on each rod proportional to their length. Ground structure methods may produce solutions where thickened

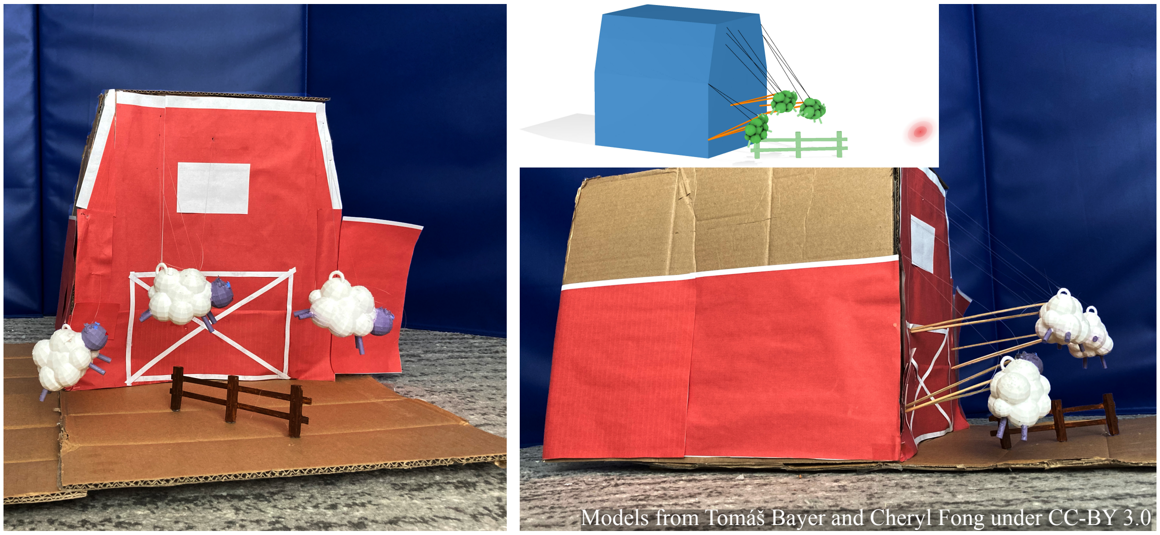

Figure 2.17: By maintaining separate graphs for rods and wires, we can use differing visibility weights and yield stresses based on what materials are going to be used in fabrication. Our system can wisely select which edges should be wires vs rods.

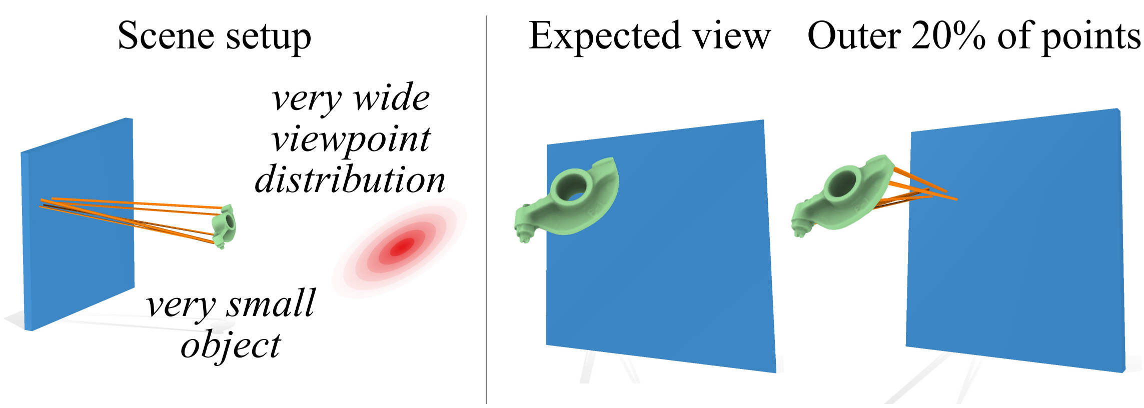

Figure 2.18: In the case of a very wide viewpoint distribution and a small or thin object, there will most likely be viewpoints from which the supports are visible. The rightmost figure shows the scene from a viewpoint on the outer 20% of the distribution.

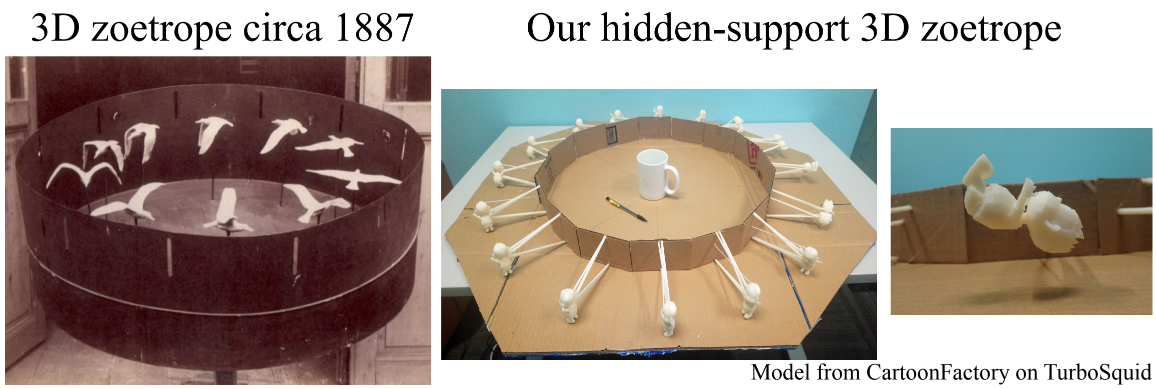

Figure 2.19: 3D zoetropes are an old idea, but hiding supports for flying objects is still challenging. We incorporate the centripetal force due to spinning and hide supports behind a backflipping boy zoetrope.

rods intersect; ours is no exception. Edges which nearly overlap with each other appear in the original ground structure and therefore may be selected as rods in the solution. However, this has not caused any fabrication problems in practice. Previous methods have considered penalty terms or post-processing to deal with intersecting (e.g., [69]). Wire-wire intersections are extremely unlikely due to the very thin nature of wires. Our visibility model considers direct line of sight, but not other cues such as reflections or shadows. Transparency of objects is not accounted for. Depending on the setup of the scene, there may not be a solution invisible to every viewpoint (e.g., Figure 2.18). Since we model physical validity as a hard constraint, we are still able to find a solution, albeit a visible one.

The precise solution depends on the initial ground structure. In general, denser ground structures produce higher quality solutions — both in terms of total structure volume and hidden-ness — with diminishing returns. Rod areas are directly proportional to stress limits, so acurate fabrication relies on accurate (or at least conservative) material measurement.

Our algorithm assumes that the input is a well-crafted scene to begin with and leaves it perfectly as inputted. The creative design process for these scenes is itself non-trivial. In the future, we are interested in pursuing an interactive design tool which would provide hints to increase occlusion by applying simple transformations (translations, rotations and scales) to the objects in the scene or even provide automatic layout optimizations given the objects and the viewpoints.

In practice assembling the rod structures with found materials can be difficult due to the blue-noise sampling on object surfaces, which sometimes places the object-rod connections at inconvenient points on the objects, requiring ad-hoc addition of hooks for hanging wires or carving of holes for inserting rods. Other future work would be to generate assembly instructions along with each solution structure.

We model the problem of hidden supports as an efficient linear program that leverages fast ray-casting from computer graphics. We see an exciting future in combining techniques from rendering and geometry processing with structural optimization in architecture and engineering. We hope this combination of appearance-driven design will be beneficial to scientific and artistic endeavours.

2.6 Appendix: Matrix Form

Solvers like MOSEK [3] expect the problem to be provided in matrix form. We spell out the coefficients of the relevant sparse matrices implementing the linear program in Eq. 2.8.

For our pruned ground structure with m candidate edges connecting N vertices, introduce a unit-less sparse matrix $\mathbf{C} \in \mathbb{R}^{3N\times m}$ where:

$$\mathbf{C}_{jl} = \begin{cases} \hat{\mathbf{t}}_{ij} & \text{if rod } ij \text{ points toward } \mathbf{x}_j \\ -\hat{\mathbf{t}}_{ij} & \text{if rod } ij \text{ points away from } \mathbf{x}_j \\ 0 & \text{otherwise.} \end{cases} \tag{2.12}$$

Here, j is used to index the 3 rows that correspond to the vertex $\mathbf{x}_j$ and l is used to index the column for rod ij.

Introduce a unit-less sparse matrix $\mathbf{Q} \in \mathbb{R}^{3N\times 2m}$ where

$$\mathbf{Q}_{jl} = \begin{cases} \hat{\mathbf{N}}_{ij} & \text{if rod } ij \text{ points toward } \mathbf{x}_j \\ -\hat{\mathbf{N}}_{ij} & \text{if rod } ij \text{ points away from } \mathbf{x}_j \\ 0 & \text{otherwise.} \end{cases} \tag{2.13}$$

Introduce a sparse unit-less selection matrix $\mathbf{S} \in \mathbb{R}^{3K\times 3N}$, where

$$\mathbf{S}_{kj} = \begin{cases} I_3 & \text{if vertex } \mathbf{x}_j \text{ lies on object } k \\ 0 & \text{otherwise} \end{cases} \tag{2.14}$$

Introduce a sparse cross-product matrix $\mathbf{D} \in \mathbb{R}^{3K\times 3N}$ with units meters, where

$$\mathbf{D}_{kj} = \begin{cases} [\mathbf{x}_j - \overline{\mathbf{x}}_k]_\times & \text{if vertex } \mathbf{x}_j \text{ lies on object } k \\ 0 & \text{otherwise} \end{cases} \tag{2.15}$$

where

$$[\mathbf{d}]_\times = \begin{bmatrix} 0 & -\mathbf{d}_3 & \mathbf{d}_2 \\ \mathbf{d}_3 & 0 & -\mathbf{d}_1 \\ -\mathbf{d}_2 & \mathbf{d}_1 & 0 \end{bmatrix} \in \mathbb{R}^{3\times 3} \tag{2.16}$$

Finally, the full linear program in matrix form may be written

$$\min_{\mathbf{a},\mathbf{c},\mathbf{q}} \quad (\ell + \lambda g')^\top \mathbf{a} \tag{2.17}$$

$$\text{subject to} \quad \begin{bmatrix} \mathbf{0} & \mathbf{SC} & \mathbf{SQ} \\ \mathbf{0} & \mathbf{DC} & \mathbf{DQ} \end{bmatrix} \begin{bmatrix} \mathbf{a} \\ \mathbf{c} \\ \mathbf{q} \end{bmatrix} = \begin{bmatrix} \mathbf{m} \otimes \mathbf{g} \\ 0 \end{bmatrix} \tag{2.18}$$

$$\text{and} \quad -\sigma_t \mathbf{a}_l \leq \mathbf{c}_l \leq \mathbf{a}_l \sigma_c, \forall l \tag{2.19}$$

$$\text{and} \quad -\sigma_s \mathbf{a}_l \leq \mathbf{q}_l \leq \mathbf{a}_l \sigma_s, \forall l. \tag{2.20}$$

where $g'_l = g_l^2$, $\forall l$ since the squared visibility is linear in the cross-sectional areas of each rod.

$\mathbf{m} \otimes \mathbf{g}$ denotes the Kronecker product of the $\mathbf{m} \in \mathbb{R}^K$ stacked vector of object masses and the $\mathbf{g} \in \mathbb{R}^{3\times 1}$ gravity vector.

CHAPTER 3

Zoetropes

3.1 Introduction: Expanding the Storytelling Capabilities of Zoetropes

Zoetropes are an old but powerful animation device defined by a sequence of images or sculptures around a spinning cylinder. They can be seen as not only a precursor to traditional film, but also as an effective, however limited, storytelling medium. A limitation of this format is the length and type of story that is possible to tell within its boundaries. Most current zoetropes are quite small, showing a simple motion of a single character that repeats over and over. We claim that this doesn't have to be the case and attempt to push the medium to tell more complex and interesting stories.

Public demand for immersive 3D experiences is evident from the success of films like *Avatar (2009)*; in fact, 93% of viewers chose to see the re-release 13 years later in 3D [23]. One of the first widely popular films to offer various 3D formats, Avatar changed the film industry and shifted the paradigm for how movies are priced and produced, with filmmakers increasingly investing in 3D cameras and technologies like IMAX and Dolby 3D. Technologies like augmented reality (AR) and virtual reality (VR) also aim to offer 3D graphics through wearable head-mounted displays. AR generally relies on glasses through which viewers see the real world with superimposed digital overlays, whereas VR usually involves a combination of a headset consisting of a screen very close to the viewers eyes and tracking of the viewer's head in real 3D space which places them in a 3D virtual environment.

The concept of 3D zoetropes dates back to 1887. Unlike 2D zoetropes' direct link to the movies we watch today, 3D zoetropes have not translated to the techniques used in "3D" film. These 3D displays are ultimately projected back into 2D; because of this VR also suffers from a lack of parallax and other visual artifacts which can perceptually affect the viewer's experience. However, zoetropes present a unique, truly 3D physical display that offers interactivity in ways that digital displays cannot. We present a modular 3D zoetrope that challenges the idea of a 3D film. Our 16-frame 3D zoetrope contains full scenes – each frame a 3D printed diorama in a box – that tell a short

Figure 3.1: Photo of the 5,200 year old Iranian vase by Michał Sałaban, Trace of a photo of the reproduction presented together with the vase in National Museum of Iran.

story. The stories we can tell are about 1-2 seconds long at frame rates of 8-16 frames per second. Our zoetrope is also constrained by its cyclical nature, but this chapter focuses on enhancing the narrative experience within these constraints by introducing user interaction and revealing layers of the story with audio.

3.2 History: From Early Animation Devices to Modern Media

Zoetropes are an optical toy invented in the 19th century and come from a long line of earlier animation devices. In fact, the first "animation" may have actually come from a 5,200-year-old bowl in Iran's Burnt City, pictured in Figure 3.1. It depicts five sequential images of a goat and are placed around the bowl in zoetrope fashion.

In 1833, Belgian mathematician introduced the phenakistiscope, a predecessor of the zoetrope which served as a proof of concept of the persistence of vision. A phenakistiscope device consisted of a disc on a handle with slits arranged radially around the edges of the disc. Inside the disc were images of the animation, as seen in Figure 3.2. When the viewer held the handle, spun the disc, and looked in the mirror through the slits, they would see the images as an animated sequence. It was used as both a toy and a scientific device, as it was being used to study the effects of pulsing lights [135].

Figure 3.2: An example of the phenakistiscope.

The zoetrope was patented in 1867 by William Ensign Lincoln. The zoetrope had a strip of photos or drawings placed around the inside of a cylinder. Above the strip was a series of slits cut out across from each frame, similar to the phenakistiscope. When the cylinder was spun at an appropriate rate, the viewer could look through the slits directly and see the animation.

The praxinoscope was invented in France in 1877 by Charles-Émile Reynaud. In contrast to the zoetrope, which used slits to view the animation through, the praxinoscope instead had a cylinder of angled mirrors around the inside of the spinning drum. The viewer could see the animation by looking at the mirrors as they spun, each frame appearing stationary in position.

In 1888, the kinetoscope was first described by Thomas Edison. Between 1889 and 1892, William Kennedy Laurie Dickson at Edison's lab worked on building it. It was a movie viewing device to be used by one person at a time where the viewer would go up to the box and look down into it through a lens. It stored a long strip of film that passed quickly between the lens and a light, and ran at 46 frames per second – way higher frame rate than the zoetrope [15].

Perhaps the most popular image associated with zoetropes is *The Horse in Motion*, a series of photos taken of a horse running, the results of which are shown in Figure 3.3. Eadweard Muybridge set up cameras in a row to be triggered by the horse as it ran by. He then went on to invent and popularise the zoopraxiscope – based on the zoetrope – and brought it to public audiences in what might be the first ever movie "theatres," where viewers could pay to see a short periodic animation of animal locomotion. The zoopraxiscope, invented in 1879, projected the images laid around the outside of a disk as it spun in front of a light. Again, in 1888 Reynaud developed the Théâtre Optique, which was an improved version of the praxinoscope that could project the images onto a screen [43]. These devices are the direct predecessors to the film projectors we use today.

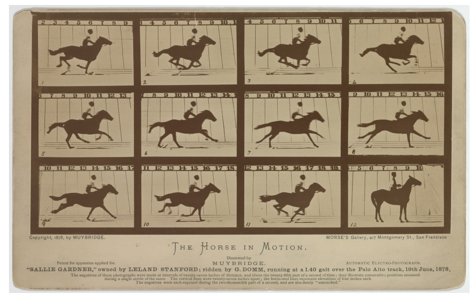

Figure 3.3: The Horse in Motion [99] series of photos by photographer Eadweard Muybridge.

When zoetropes were first invented, artists would use a small slit or window to act as a sort of shutter. Viewers would peek through the slit and see the image for a split second of time. The slits help to trick the eyes into seeing motion from the series of still pictures, by blocking the empty spaces between the images [85]. Zoetropes illustrate the concept of apparent motion by converting a series of still pictures into a display of continuous motion. This phenomenon is the basis of any form of "moving" visual media, such as film. Due to the way the human brain processes visual cues sent from the eyes, flashing a series of images at any rate above 8-12 frames per second will be perceived as motion as opposed to the still images truly present.

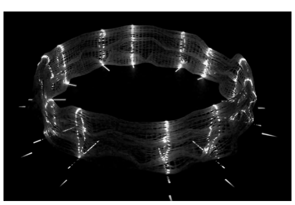

Figure 3.4: CROSSING #3 on display by Goto [48].

Though they may seem like outdated contraptions, zoetropes are still made and loved by people today. Big animation studios like Pixar and Studio Ghibli have 3D zoetropes on display with figures of their beloved characters positioned around the disc. Zoetropes do not always have to be cylindrical; they can also be linear. In this case the frames are not constrained to be in a loop, but instead laid out from beginning to end in a line. A fun example is the Masstransiscope, which is located on a New York City subway line. Instead of the viewer being fixed and spinning the zoetrope, the zoetrope is fixed to the inside of the subway tunnel and the viewers on the train looking at it through the windows as they travel act as the shutter which creates the motion [57].

Artists like Akinori Goto and Gregory Barsamian currently make zoetropes and exhibit them throughout the world [10]. Barsamian makes large spiral zoetropes that contain sculptures around the metal frame. The tall structure spins and uses a strobe light to replace the traditional zoetrope's slits. His animations usually consist of transforming objects from one state to another in their journey from the top of the spiral to the bottom. Goto makes 3D printed circular meshes made from nylon that look, at first, like a confusing and meaningless object. However, he uses clever light projection methods where he shines a continuous strip of light onto the mesh. As it spins, the "frame" under the light is seen as an outline of the shape at that point (see Figure 3.4). The light does not strobe and therefore the animation appears more smooth compared to the choppiness of the other methods including the strobe light, slits, and mirrors of the inventions listed here.

One thing all these devices lack is interaction between the viewer and the animation. If just watching these animations in person is so magical, imagine how engaging it would be to put the power of experiencing the story in the hands of the viewer.

3.3 Related Work: Why Audio and why zoetropes?

3.3.1 3D Displays

3D animated displays have long been a challenge to create effectively, with various technologies presenting their own limitations. VR systems, for example, often suffer from issues such as vergence-accommodation conflicts, which can cause discomfort for users [78]. Stereoscopic projection suffers from limited depth perception as well, and polarizing glasses – which block light from each eye to create its 3D effect – can lead to reduced brightness and colour quality of the projected image [64]. Additionally, both VR and stereoscopic projection require the viewer to wear a headset or glasses, which can be uncomfortable especially for viewers that already wear prescription glasses. Stereoscopic monitors frequently struggle with poor focus and visual "crosstalk", where images meant for one eye are partially visible to the other, leading to subpar visual experiences [90, 141].

Volumetric 3D displays offer a higher quality of depth perception without the need for glasses. These types of displays either use reflection or emission of light through various lenses or screens, for example shining light through a spinning mirror [70], projecting images through a stack of liquid crystal shutters [131] or onto a thin layer of fog [117], or emitting spots of light in 3D space using modulated laser pulses [103].

Another such type of volumetric display is a light field display, which provide a holographic 3D image by displaying multiple views of the scene at the same time, using a high resolution screen and a lenticular lens array on top which makes them expensive to produce and hard to make content for. They are also forced to be relatively low resolution because the lenses over groups of pixels have to reflect light to many different viewing angles, and even so, their 3D effect is limited to a small range of field of views [139].

Work has been done to interact with these volumetric displays such as bringing the 3D image into focus at any position around the 360 degree display by tracking a viewer's head [70] or even to use direct gestures to manipulate the holographic 3D objects by tracking the user's hands around the display's enclosure [49].

VR and AR displays are both backlit, while stereoscopic and volumetric 3D displays project light onto the viewing surface/volume. These technologies share a common drawback: they emit light out toward the viewer without the capacity to interact with incoming light, limiting their potential for interactive storytelling using light.

Meanwhile, mechanically animated 3D displays, such as animatronic robots with articulated characters, require a high level of technical skill and are difficult to produce at scale. The complexity of the story is limited by physical articulation of joints and expensive robot parts and electronics.

3.3.2 Zoetrope Advancements

Zoetropes, in contrast to animatronics, instead of having a single frame with 3D moving characters, have many frames of still 3D characters. Each frame does not have to move realistically, eliminating the need for skilled engineers and artists to create a convincing story. With the popularity and availability of 3D printers, creating complex 3D scenes in the physical world is now easier than ever.

Expanding the capabilities of zoetropes is a current endeavour. Smoot et al. created a zoetrope with strobing lights which can project real-time, non-repetitive animations onto solid figurines or holograms, de-limiting the amount of unique frames which can be used on a single zoetrope [127]. One version of their device uses 8 ping pong balls with drawn cartoon faces with different mouth positions as frames. A viewer can speak into a microphone and watch the ping pong ball animate to mimic their mouth movement by spinning at a high enough pace and strobing the frame with the mouth open at a position proportional to the audio. This innovation allows for infinitely long and interactive animations on a zoetrope, but is limited to a simple shape. Yokota et al. has proposed a multi-layer 3D zoetrope which can display two independent animations concurrently, by way of concentric turntables, two strobe lights and a semitransparent mirror [146]. This doubles the possible frames of an animation using 3D models, and allows for interesting superimposing of animations, but ultimately there is still a relatively low limit on frames and no improvement on interaction. The Eigen zoetrope works by strobing a light at different intensities on 8 different rapidly spinning images to superimpose a unique animation, effectively blending the basis images into 18 new frames [75]. Like the Eigen zoetrope, the ZoeMatrope [98] also linearly blends frames, but in this case, to simulate realistic material properties. It employs a spinning disk with 6 differently coloured frames of the same object to act as the base materials. By varying the intensity and emission time of the strobe light on a combination of the frames, they are able to compose both diffuse and specular material parameters. These techniques focus on expanding the zoetrope by way of un-restricting the physical number of unique frames that can be continuously displayed.

Our previous demo proposed specific improvement on engagement through interaction by putting a strobing flashlight in the hands of the viewer [79]. This gives the viewer freedom to illuminate specific parts of a zoetrope's animation, creating an opportunity for a new mechanism to expand on a zoetrope's story by having changes in illuminated areas create some sort of response.

3.3.3 Audio in Storytelling

As soon as technology developed to make silent films into "talkies", it became widespread [29]. It is well known that audio enhances visuals in a variety of ways. A lot of information can be encoded into audio: a dialogue, music, ambient sounds for example, can convey mood, give a story context, or set the scene. Sound design in film has a profound emotional impact and can change or amplify the way a scene is perceived [47]. Sound has also been shown to influence a viewer's attention positively of an otherwise not very stimulating visual cue [96]. Sound can make a storytelling experience more immersive; Privitera et al. explored the effect that sound has on enhancing visitor experiences at cultural heritage sites like museums, and studied how audio increases engagement by making the story personal to the viewer [113].

Our interactive zoetrope allows the viewer to shine a strobing flashlight into the scene, triggering audio cues that enrich the visual story. With 16 frames of animation, the zoetrope offers an opportunity for incremental storytelling – where plot elements and emotional textures are revealed gradually as the viewer explores the scene. Additionally, the same visual sequences can be used to tell multiple stories by altering the accompanying audio, further expanding the potential narrative depth. By engaging the viewer's imagination and emotions, this approach helps extend the domain of stories that can be told using a zoetrope, making it a more dynamic and engaging medium.

3.4 Method: Enhancing Zoetropes with Light and Sound

We introduce what we call the Audiotrope: a light-reactive 3D zoetrope display with audio components. In the following section we describe the hardware that makes up the Audiotrope, the 3D-printed animation process, how we enable interaction with the light, the pre-processing steps needed to prepare the animation, and finally how to run the Audiotrope and watch the 3D animation.

3.4.1 Anatomy of the Audiotrope

Figure 3.5 shows a labeled image of the important pieces of our zoetrope as we describe in this section. The Audiotrope's 16 frames are fastened on the outside of a bike wheel mounted in an aluminum extrusion structure. The 30 x 30 mm aluminum extrusion rails are lightweight but sturdy and easy to assemble with corner brackets, lock washers, and bolts. The bike wheel sits on a 14 mm diameter steel rod and spins thanks to 5.5 mm ball bearings on both sides of the axle. It is mounted on the 0.45 x 1.2 x 1.8 m extrusion cage at a height of 1 m from the ground.

Figure 3.8 shows a flow of data through the Audiotrope's hardware system. A Nema 23 bipolar stepper motor driven by a DM542T motor driver rotates the wheel using a bike chain connected by a sprocket (mounted to the motor shaft using a 3D part we designed in Fusion 360 [5]) and the gear from the bike wheel at a 17:9 gear reduction ratio. The Audiotrope turns on by pressing a big red push button, which can be seen in Figure 3.5 (a). The button then lights up and an Arduino Mega (the first microcontroller) begins the process of ramping the motor up to full speed. We chose to use an Arduino Mega because we needed a large number of pins to implement the motor control loop. Like all stepper motors, we control the 3Nm 4.2A motor by specifying the direction and step. By default the motor steps 200 times per revolution. This is acceptable for a regular zoetrope, but we need more control over the rotation when we add the audio components. We set the driver to microstep at 6400 steps per revolution. The motor driver also has a braking feature, where we can disable the drive when the zoetrope is off for safety and stability.

The wheel with all the scene holders and scenes clipped in can be very heavy and may not be uniform weight throughout so it is important to make sure the weights are counterbalanced using the same techniques that motorcycle mechanics use. The weights can be put on any spoke of the wheel and tightened with a hex key. The weights can also be slid up and down the spoke to be closer or further from the axis of rotation, giving more control over the balancing. There are four different weight options, shown in Figure 3.5 (f). Motorcycle spokes are thicker in diameter than regular bike spokes, so we 3D-printed small inserts for the weights to be adapted to fit around the spokes of our bike wheel.

Another crucial element of the zoetrope's smooth spinning is the bike chain tensioner that is mounted to the aluminum extrusion and uses a spring to prevent the chain from becoming too loose

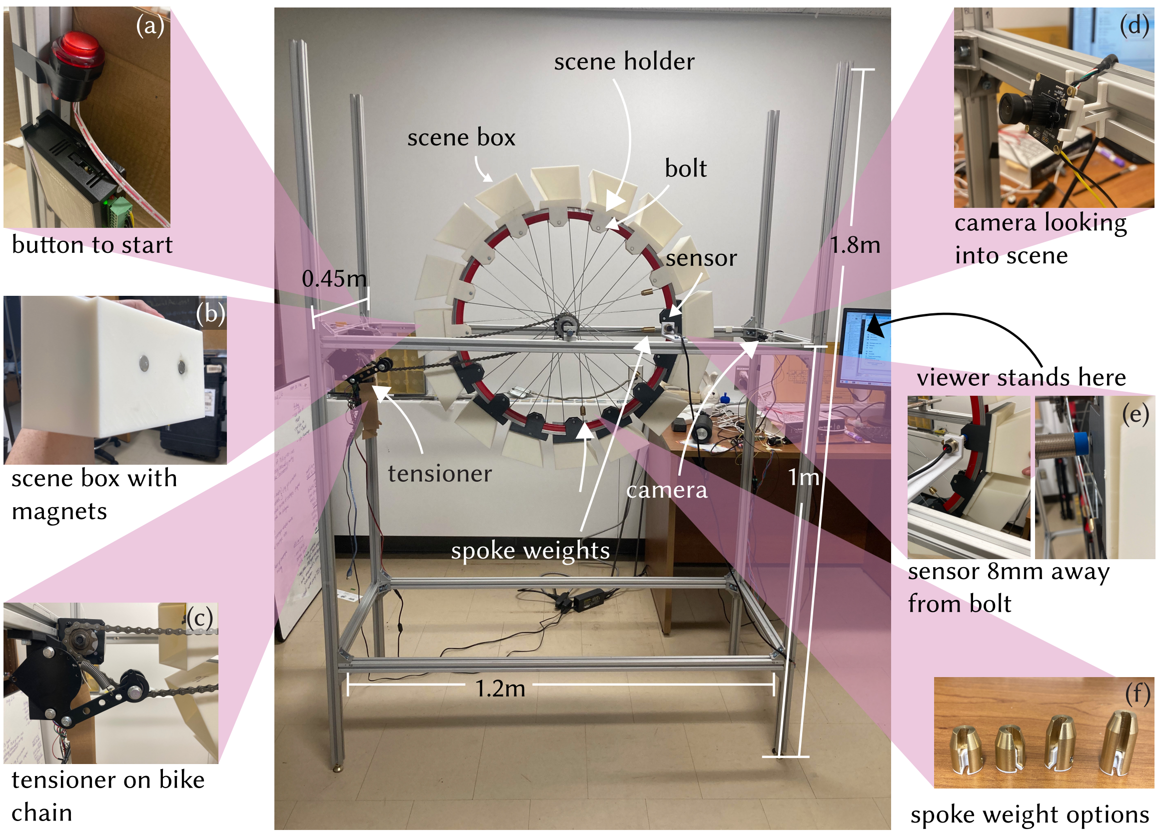

Figure 3.5: A photo of the inner workings of the zoetrope: the start button to ramp the motor up to full speed (a), a scene box with magnets to be clipped into the scene holder (b), the tensioner on the bike chain to maintain smooth rotation (c), the camera pointing from where the viewer stands into the current frame's scene box (d), the proximity sensor for triggering the strobe light when it detects a bolt (e), and the spoke weights for wheel balancing.

or too tight and keep an even tension, as seen in Figure 3.5 (c).

As the wheel with its attached scene boxes spins and the bolts holding them to the wheel pass an inductive proximity sensor (Model LJ18A3-8-Z/BX) mounted onto the extrusion via a custom 3D-printed mount, the digital signal read by an Arduino Uno (the second microcontroller) triggers the flashlight to strobe and the Innomaker U20CAM-9281M camera to capture an image. This control loop is separate from the one driving the motor and gets its input from the sensor triggered by the wheel spin. The sensor has a detecting distance of 8 mm, which means it has to be mounted in such a way where the bolts on the scene holders will be close enough to be detected but not close enough to hit the sensor as it passes by. See Figure 3.5 (e) for a close-up of the sensor. The 3 wires of the sensor plug into ground, 5V power output, and a digital pin with internal pullup on the Arduino Uno and allow us to read a digital signal when a metal piece comes within 8 mm of the end of the sensor.

To strobe the light, we designed a small PCB to modify the Pocketman mini flashlights, as seen in Figure 3.6. Normally the flashlights are battery-operated and turned on with a button to shine continuously. As stated, zoetropes create apparent motion using a quickly flashing light, so our flashlight needs to be easily triggerable by the Arduino to turn on and off for a short period of time. We opened the store-bought flashlight, took out its LED driver board, replaced it with our controller, and then put the flashlight back together. The PCB is soldered to the LED in the flashlight, as seen in Figure 3.6, and uses a stereo 3.5 mm plug to connect to a 3-pin screw terminal, which allows connection of ground, 5V power, and trigger into the Arduino. The length of the flash is short to reduce blurring effects caused by the motion of the wheel. The flash time can be modified to be shorter or longer, with the caveat that a longer flash time is easier on the eyes but will make the frames appear more blurry.

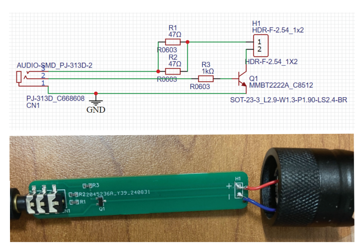

Figure 3.6: Flashlight controller schematic and PCB.

The camera is mounted – again, using 3D-printed parts – on the same rail that the bike wheel and sensor are mounted on, pointing into the scene just below the viewer's eyes, as seen in Figure 3.5 (d). Two wires soldered to the camera go to the Arduino Uno for ground and trigger to allow the camera to quickly capture an image in the short amount of time that the light is on. Once the camera takes an image, the data is sent via USB connection to a Raspberry Pi 4B, which runs a Python program to analyse the images and isolate the flashlight beam to see where the viewer is looking in the scene. Learn more about the image analysis in Section 3.4.3 If the viewer is looking at a specific frame at a specific object in the scene, audio cues can be triggered to enhance the story and make the viewing experience more interactive and fun. Audio files associated with objects in the scene are played through a JBL speaker that is connected via 3.5 mm audio jack. It can also be connected to wired headphones which could allow for a more personal listening experience. In the following sections we describe in more detail how the the Audiotrope works and the process of creating an animation to be played.

3.4.2 Animating for a Zoetrope

All animations made for our Audiotrope were made using standard animation software, in our case, Blender [26]. Typically digital animators create their animations at 30 frames per second. Zoetropes

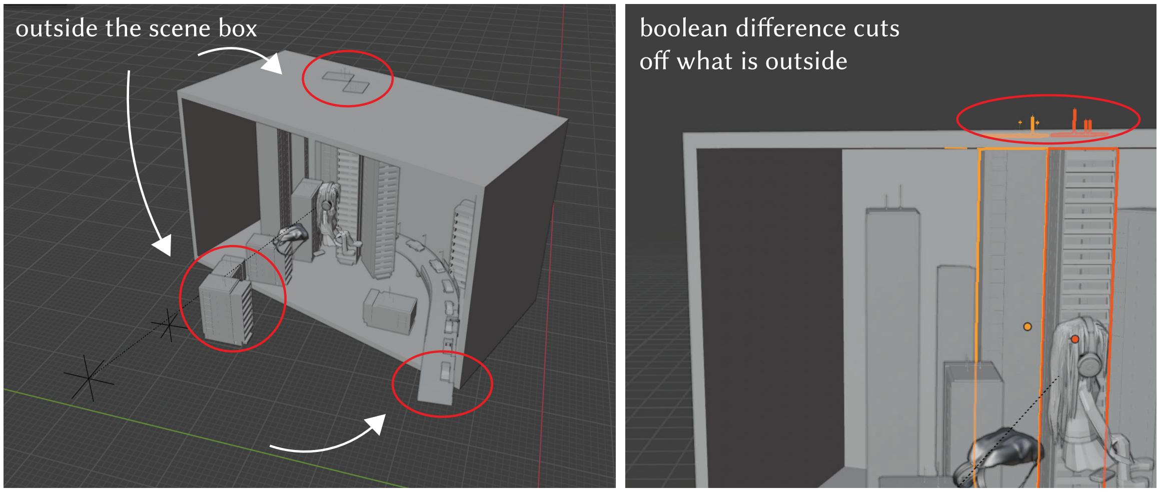

Figure 3.7: By differencing the building meshes with the scene box geometry, what is outside the box is separated, deleted, and therefore not 3D-printed.

run at a much lower frame rate – around 8 frames per second. We choose to use fused deposition modeling (FDM) 3D-printing as our mode of fabrication because of the speed and availability of 3D-printers. Sixteen frames need to be 3D-printed to fit around the bike wheel. The animation is regularly sampled, taking every third or fourth frame in the original short animation. We use motion-capture animated rigged characters from Mixamo [1], which lessens the burden of keyframing characters' motion by hand and can be downloaded in FBX format at 30 frames per second.

The animation will eventually have to be fabricated, which requires careful selection of models with geometry that can be easily 3D-printed. Digital modelers and animators usually have other goals when they do their jobs. The focus of their task is either to make the models look good – this means that the models can have non-manifold geometry, overlapping shapes, and thin surface meshes with no volume – or to make them fast to render and animate so they can be used in video game engines – this means that the geometry can be low-poly style and have fewer details. In both cases, consideration for 3D-printing is not necessarily a priority.

3D-printing software, called slicers, generate commands that tell the extruder on the printer what path to go in order to make the desired shape. Slicers often expect meshes with "good" geometry as input. The slicer attempts to create the paths which approximate the shape best, and with features like non-manifold edges, the algorithm struggles and sometimes fails to slice the object. To mitigate these issues, trying to find objects that have good geometry for printing at the start is crucial. When this is not possible, various Blender shape modifiers like Solidify, Decimate, Subdivide, and Remesh can be used to get the models in good enough shape to be represented correctly in the slicer. This can be a trade-off between detail and quality of the mesh, and the need to be sliced quickly and correctly.

The 3D-printed diorama needs to be printed in the proper scale to fit into the zoetrope. The scene boxes are 9 x 16 cm in height and width and can be a variety of depths depending on the scene. For example, the buildings of a vast cityscape or the highway through the city may come out of the dimensions of the box. To fix this problem, boolean difference operations in Blender cut geometry that is outside the bounds of the box, as seen in Figure 3.7.

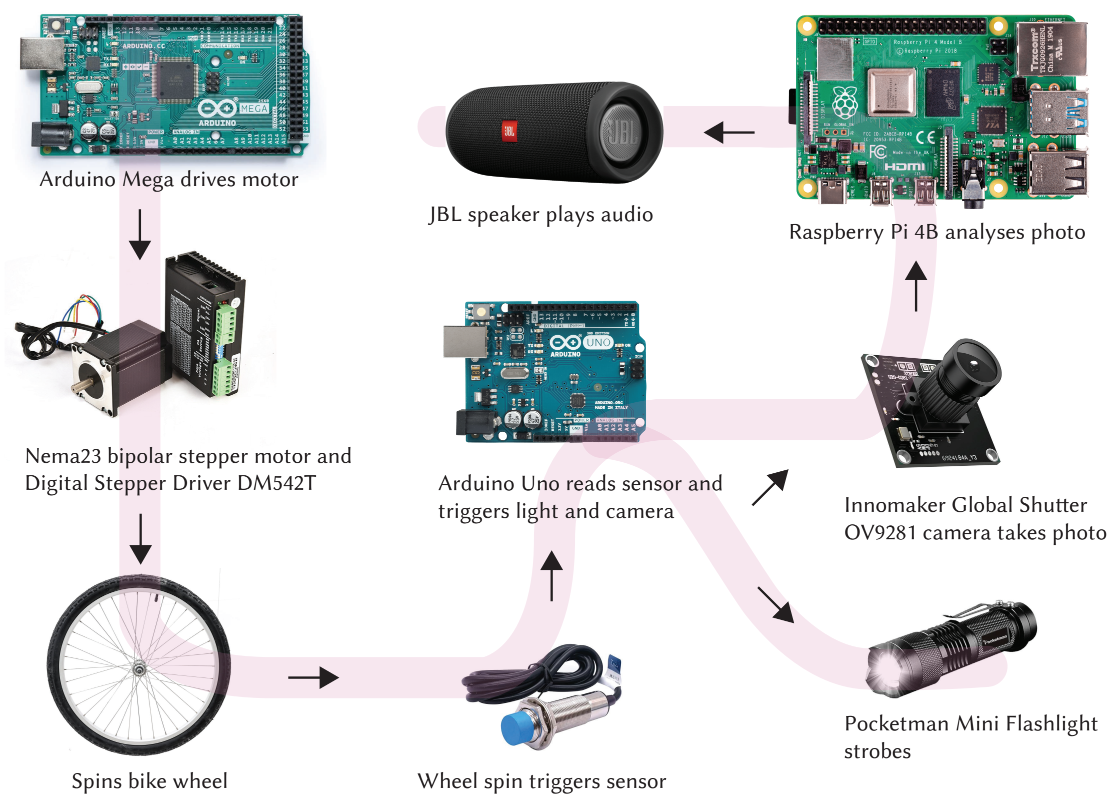

Figure 3.8: The flow of data through the Audiotrope system with hardware labeled. Items are not to scale.

Another consideration during the printing process is keeping track of which frames have been printed and which still need to be done. To ease the printing process, we run a python script in Blender to carve out the frame number of the geometry so it is labeled after it is printed (see Figure 3.9).

After printing each frame and its scene box and assembling the scene, the artist must glue extra-strength magnets into the holes in the back of the scene box (being mindful of the polarity!!) so they can be attached to the scene holders on the bike wheel (see Figure 3.5 (b).)

Figure 3.9: The bottom of the girl's platform has her frame number carved into the bottom.

3.4.3 Interaction with Light

As we have seen, zoetropes are limited in the types of stories they can tell. In order to give a zoetrope story more depth, we introduce light-reactive audio feedback to enhance the meaning, and audience experience of the story. Because zoetropes use strobing light, we can take advantage of that light and introduce it as a user tool for viewing the animation. The strobing flashlight is triggered to flash for 500 microseconds when each frame passes the viewer.

Our Audiotrope needs to be able to both detect where in the scene the viewer is pointing the flashlight and trigger audio to play if the viewer is looking at what we call *Action Objects*. Two pieces of information needed in order to do this are the following: what frame the viewer is seeing, and where in the scene they are looking.

In order to know what the viewer is focusing on, our system needs to know what frame the viewer is currently seeing between 1 and 16, since object locations can vary from frame to frame in any animation.

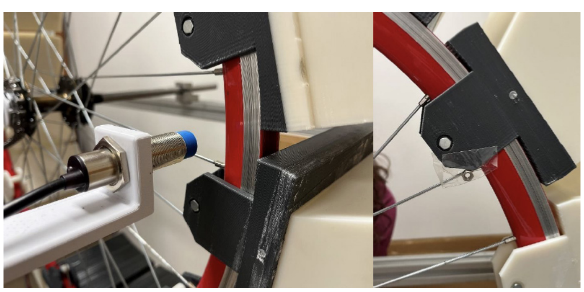

Figure 3.10: The last frame of the zoetrope has a small additional metal nut, indicating that we can restart the count from 1 to 16.

As mentioned previously each scene holder is fastened to the zoetrope with a metal bolt which can be detected by a changing magnetic field when the sensor is close to the bolt. Only the last frame contains an additional metallic nut to be detected (see Figure 3.10). We keep track of how long the last digital signal received from the sensor was, which is usually between 40 and 100 milliseconds once the wheel is spinning at the proper speed for the animation. The elapsed time is not exact because of imperfections of the wheel alignment, which affects how the sensor can detect the bolt passing by. Since the bolts on each frame are spaced regularly, the elapsed time between regular frame detections will be consistent, and the last frame will have an additional detection with a much shorter elapsed time. The shorter elapsed time is used to reset the current frame counter back to 1. The current frame number is sent serially through USB to the Raspberry Pi.

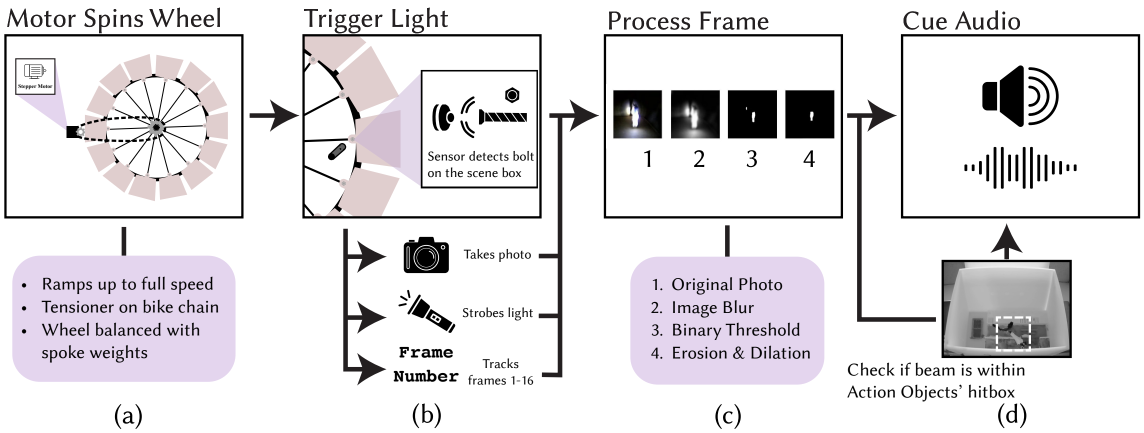

Figure 3.11: A summative diagram of our zoetrope's control loop. The wheel's spinning (a) triggers a light and camera (b) and the image goes through an image processing stage (c) so we can check if the viewer is looking at an object. If the viewer is seeing an Action Object, audio is played (d).

For object tracking, we use a UVC (Universal Video Class) compliant USB camera facing into the scene, just below the viewer's eyes, which is triggered through hardware when the sensor detects the bolt. The light flashing needs to be synced almost perfectly to the camera shutter so that the image taken by the camera contains the beam of light from the flashlight. Once the frame is captured, simple image processing operations – Gaussian blur, thresholding, then erosion and dilation – are performed on the image to isolate the beam of light, as shown in Figure 3.11 (c). The isolated beam of light in the image may be on one of the action objects, and if it is, audio needs to play. Before the animation plays, for each action object in each frame, it is necessary to know the x, y coordinates of the "hitbox", a rectangular boundary containing the object, so that its corresponding audio can be triggered.

3.4.4 Calibration Process

The calibration process that happens before an animation can be played on the audiotrope asks for vital information about the scene, like which objects trigger audio, in what frame, and what audio file is associated with it.

We use template matching, an algorithm used for searching for a smaller template image's location inside a larger source image [18]. Template images consist of the action object within a particular frame, while the source is that entire frame. The algorithm essentially slides the template pixel-by-pixel over the source image and for every pixel of the template that overlaps with the source, a metric is computed and stored in a result matrix. The metric defines the distance, or inversely, the similarity between the template and the patch of the source image [52]. The result matrix describes how good of a match each portion of the source image is to the template image. This matrix represents the top-left vertex of the best match as the brightest point, which allows for coordinate extraction. The coordinates of the bottom-right vertex are then obtained by adding the width and height of the template image to the top-left coordinates, and both sets of coordinates can be stored and used to define the location of a particular action object within a particular frame. We use OpenCV's Python bindings to compute the hitbox of the action objects [13].

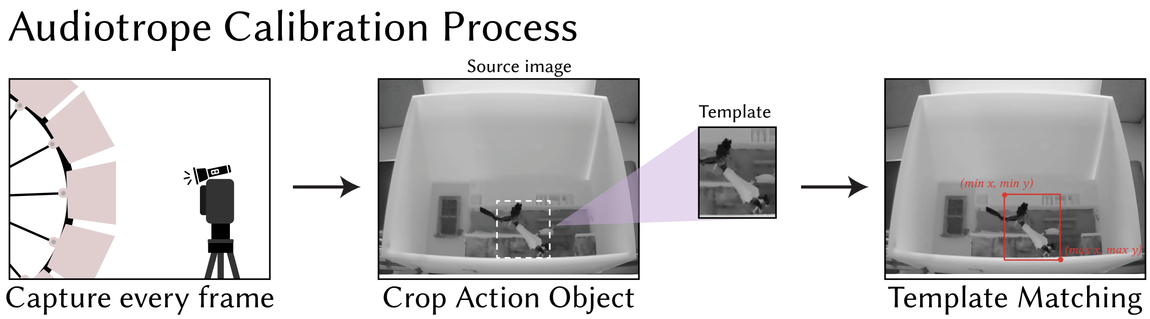

Figure 3.12: Calibration process diagram. Running OpenCV's template matching algorithm on frame 9 on the Pancake Flip animation shows that the girl falling template image is correctly identified within the whole scene source image.

The first step in the calibration process is to take an image of the whole scene in each frame and save it. Then, Action Objects must be cropped out of each frame, acting as the templates for the source image of the whole scene. Once this cropping is done for every object in a single frame, it has to be repeated for the other 15 frames. This process can be tedious and has potential improvements to be made in the future; however, since the animations are only 16 frames and the number of action items is assumed to be small, it is a reasonable solution and only has to be done once. For each frame and each object, the template matching function is run and the objects' rectangular coordinates within the scene are recorded.

After the imaging phase is finished, the next step is to tell the system which audio files should be played when a given action item is looked at in a given frame. In order to associate object and frame information with an audio file, the zoetrope artist goes through an interactive process of selecting directories on their computer containing the cropped template images and the whole scene source images. The directories must be organized in such a way that our system can read them (see Figure 3.13).

There are two root level directories: `images` and `audio`. For every action object, there should be a sub-directory named as the action object in the images directory containing the templates for each frame that the object is in. Similarly, there should be a matching sub-directory in the audio directory which contains options for clips that could be played during those frames when the object is being looked at.

We use our own JSON format to represent the python dictionary data structure that says which audio belongs to which object. We take in the file structure, do template matching to obtain object coordinates, and output a JSON file that can be saved to the Raspberry Pi and loaded in at the start of the animation playing process.

3.4.5 Playing the Audiotrope Animation

The animation playing process involves setting up the physical contraption, as described in Section 3.4.1, performing the calibration process in 3.4.4, and starting the main Audiotrope control loop. Assuming the zoetrope artist has animated and 3D-printed 16 frames and clipped them into the scene holders, the audiotrope is ready to be played.

The Audiotrope has to be situated in a room where the lights can be turned off so that the strobe light can do its job and create the illusion of motion. The person running the Audiotrope first has to start the wheel by pushing the ON button. This starts the motor and therefore triggers the light and camera. Separately, the camera's images are sent to a Raspberry Pi 4B, where our main script must be started.

Our main program loads in the JSON file containing image names, audio file names, and action object coordinates created during the pre-processing step. Then, it waits for an image and its frame number to be sent to it. Once the image is read by OpenCV, we perform the simple operations described earlier to isolate the flashlight beam and check whether any of the objects' hitbox coordinates contain the pixels where the light is. If an object is detected and that object appears in the current frame number being received, one of the audio file options is chosen at random and played on a speaker. We use the PyGame library available on the Raspberry Pi operating system [114] to play corresponding audio in real time as the viewer looks into the scene with the strobing flashlight.

A viewer could be shifting focus from object to object at any given time, so audio needs to be played, paused, or stopped and maintain synchronicity with the animation. The PyGame library has a `pygame.mixer` audio module that allows easy playback of multiple audio files at a time. It also provides functionality for checking if a channel is currently playing audio, allowing us to stop and start it as needed. We create mixer `Channel` objects for every audio file belonging to the whole animation. During the real-time process, after locating the object in the current frame, we play the correct channel with the audio file that is picked from the audio options available. PyGame can also control the volume of the audio being played, represented as a percentage of the volume that the speaker is set to. When the viewer changes focus from one action object to another, the current file being played is faded out over the course of 1 second to avoid abrupt stopping and starting. Because we pause the channel instead of stopping it completely, the audio file picks up where it left off once unpaused.

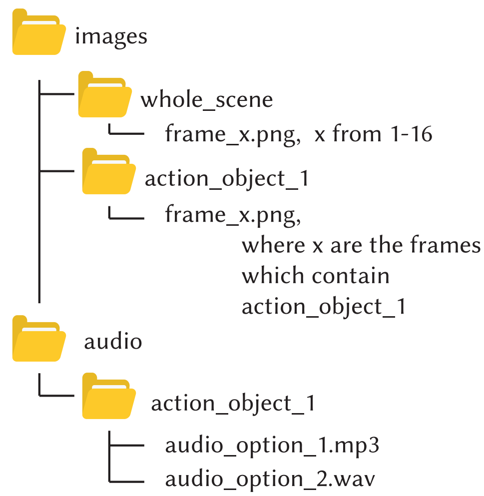

Figure 3.13: All files must be arranged so that the system understands which objects trigger which audio clips.

Now the viewer can turn off the lights, hold the flashlight, and be immersed in the story being told visually and explore the scene to gradually reveal more through audio.

3.5 Audiotrope Stories

We present examples of stories that we created with our method and propose another story idea, which could be implemented in the Audiotrope using other features.

3.5.1 Pancake Flip

A person is cooking a pancake in her kitchen. When she tries to flip the pancake, she flips it too far away and has to jump after it. At the end of the story, she falls to the floor and misses the pancake, putting her head down in disappointment. Our first animation was used to validate our Audiotrope

Figure 3.14: An entire kitchen environment fits into a scene box where a character flips a pancake.

method, as well as demonstrate the success of our hidden support generation algorithm. The entire animation is 60 frames at a rate of 30 fps, making the 16-frame zoetrope rate 8 fps.

For the Pancake Flip animation, pictured in Figure 3.14, we used a Mixamo character and rigged her using the Blender add-on Rigify. We keyframed the animation of the girl by hand using a reference video from YouTube of a person flipping a pancake. The pancake was animated using Blender's built-in rigid body physics simulator. We gave the pancake a weight and an initial force vector on the frame when the pancake leaves the pan. For every frame, we run our method from Chapter 2 to obtain a rod structure that invisibly holds up both the girl and her pancake as they fly through the air (see Figure 3.15).

In this animation, we have 3 Action Objects: the pancake, the girl when she's falling, and the girl when she's landing. The girl is falling in frames 11 and 12, so in the `images/girl_falling` directory, we have the templates cropped from the whole scene in those frames – `frame11.png` and `frame12.png`. We have two audio options in the `audio/girl_falling` directory – `thud.wav` or `clang.wav`, so either of those could be played when the viewer focuses on the girl with the light at frame 11 or 12.

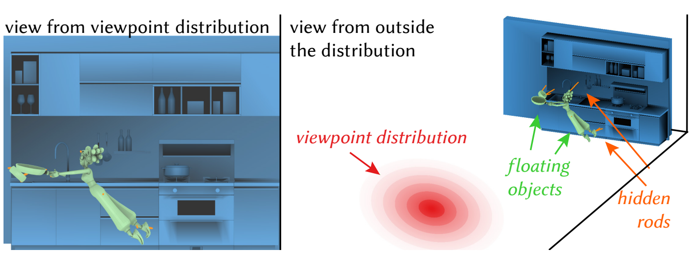

Figure 3.15: Using the method from Chapter 2, we are able to invisibly support the girl and her flipped pancake while they fly through the air.

When the girl is flying through the air, the viewer can hear her yell "not again!" or "aaaaah!". When the pancake is still in the pan at the beginning of the animation, the viewer can hear it sizzling when they focus on it.

Zoetropes are periodic in nature, and the pancake animation is not periodic. For this reason, after the 16 frames happen and she is on the ground with her pancake, she snaps back into position next to her stove at the beginning of the story, frame 1. Maybe this isn't ideal for a zoetrope story, but it does not affect the story or viewer in a negative way.

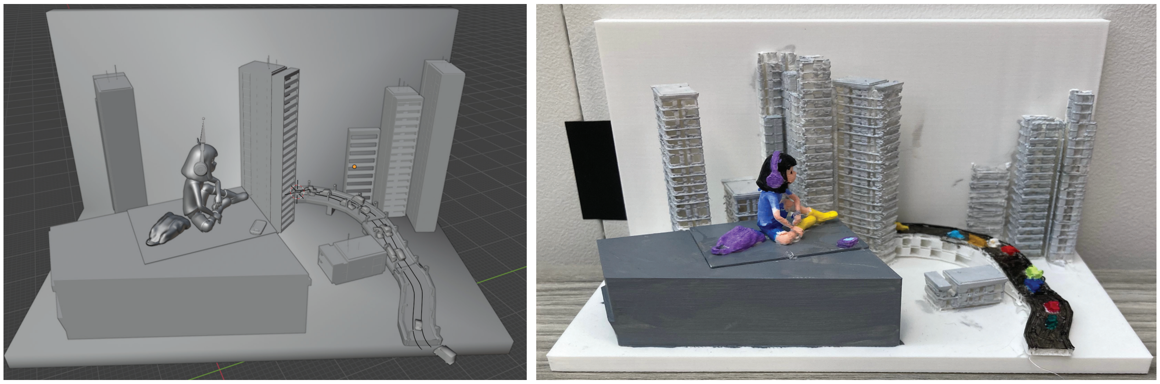

Figure 3.16: The City Girl perched on her building digitally in Blender (left) and physically printed and painted (right).

3.5.2 City Girl

In the City Girl animation, we explore the idea of using the same visuals but with multiple scenarios told through different sets of audio. The animation is 45 frames at a rate of 30 fps, making the 16-frame zoetrope rate 10-11 fps.

In this story, a high-school aged girl sits atop a building bobbing her head to the music coming through her headphones. Next to her, her phone is ringing, but she either doesn't hear it or chooses not to answer. Behind her in the buildings in the distance are different people going about their day. Most importantly, her parents are seen having an intense conversation in one of those buildings. Cars pass through the city on a highway below. In the background, an ambient city sound plays quietly.

We leave the visuals simple and intentionally vague so that the audio can tell distinct stories. The action objects in this scene are the girl, her phone, the highway, the building where her family lives, and two other buildings in the scene. When the viewer focuses the light on the girl, we can hear her music or the thoughts going through her head. Her phone rings when the viewer looks at it. Maybe after a long time of the phone ringing, we hear the voicemail left by the caller. The cars on the highway make sounds. When the viewer looks at the building where her parents are, the conversation between mom and dad can be heard. One building has a person watching a sports game on TV and another contains animals playing poker. The sounds coming from the buildings behind also have a 2D animation that plays.

We have two – but theoretically there could be many more – versions of the story. In the happy version, the girl listens to upbeat music. Her inner monologue shows that she is excited but nervous because she is expecting results from her university applications. She is waiting to see whether she got accepted into her first choice school. Meanwhile, her parents have just received a letter in the mail. It's her acceptance letter, and they begin to plan to surprise her with her favourite meal for dinner. They talk about how proud they are of her. Her sister is calling her on the phone, but the girl's music is too loud. After a while, the sister leaves a loving voicemail telling her to hurry up and come home before it gets too dark.

The sad version is told through the other set of audio clips. The girl listens to angsty music. Her inner monologue implies that she is staying out late on purpose and doesn't want to return home.

Her parents argue in the distance about what they should do and think after they find out their daughter is gay. They both disapprove but disagree on the best way to handle it. The girl's thoughts suggest that she hears her phone ring but chooses not to answer. She doesn't want to answer the phone because she thinks it's her parents calling. Her best friend is calling, telling her that she loves her and wants her to be happy, after having initially said something hurtful.

3.5.3 Monster Under the Bed

This is an animation that we created digitally but did not have time to print. A boy sits in his bedroom, swinging his legs on his bed. Seemingly, there is nothing else interesting about the scene. When the viewer is looking at it as a whole, i.e. the beam of light is wide and covers the whole scene, it is just a boy in his room. When the viewer takes a closer look at the scene with the flashlight, the beam will be smaller and more focused as it approaches the scene. As they look around the room and the light lands on the curtain in front of the window, suddenly the viewer hears the sound of a dog barking and see a dog there! Other spots in the room trigger noises and other animals are revealed. Audio could indicate that the boy's parents are telling him not to be afraid of the monsters under his bed. When the viewer peeks under the bed, a scary monster appears. So how would this be implemented?

Our zoetrope has 16 frames, but only 8 frames per second is "required" for the phenomenon of apparent motion. This means that two different animations could be interleaved in sets of 8 frames. In other words, the entire animation would be 1 second long, but the odd numbered frames could have a slightly different scene than the even numbered frames. The odd numbered frames show the boy in his bedroom with no animals present. When the beam of light is large, the zoetrope only strobes the light on the odd frames. When the beam of light becomes smaller, the script on the Raspberry Pi could detect this, along with the location in the scene the viewer is looking that we already compute, and switch to strobing the light on the even numbered frames. In this frame set, the dog behind the curtain and the monster under the bed are there. But because the light is only shining on a small area of the scene, the viewer only sees the animal they are looking at and the rest of the animals stay hidden in the dark.

Implementing this requires 3D-printing the 16 frames, and augmenting our data flow from the Arduino Uno to the Raspberry Pi to go both ways. The current frame number being shown to the viewer is sent from the Uno to the Pi. In this case we would also want to send information back to the Uno that communicates that it needs to switch to strobing the light on the even or odd frames. We would signal to switch when the image processing step on the Pi shows that the viewer is focused on an area of the scene with an action object, The rest of our system would work as is.

3.6 Future Work: Narrative Complexity & Accessibility

As stated in Section 3.4.2, geometry with undesirable characteristics can make the 3D printing of the animation more difficult than expected. The 3D printing software, hardware, and filament properties can all pose problems, which eventually lead to the resulting print not correctly representing the digital object. Manually fixing these problems takes time and ideally should not take time away from the zoetrope artist, whose priority is the animating and the storytelling. In the future, having an automated tool for improving the geometry before printing by scaling, thickening, or smoothing parts of the object, for example, would greatly improve the animation process.

Even creating the story and the digital animation itself can be difficult. Emerging artificial intelligence tools which generate visual content from text using large language models (LLMs) also promise an easier future for animation creation [115]. Although these models can only generate 2D video for now, generating animated 3D meshes is not far behind [105, 37]. Progress has been made on generating neural radiance fields (NERFs) [94] – a different 3D digital representation, but NERFs do not capture fine details and are not easily animated or 3D printed. Additionally, in combination with chatbots like ChatGPT [106], artists can collaboratively brainstorm and refine their story lines with the help of an LLM even before getting to the animating stage.

We discuss accessibility concerns that come from the strobe light. The light is on for such a short amount of time, between 200 and 1000 microseconds 8 times per second. This is extremely problematic for viewers with epilepsy or other disorders that prevent the ability to look at strobe lights. Even viewers without this problem can report discomfort from looking at strobe lights for too long [39]. The flicker fusion threshold tells us that humans stop seeing flashing lights and our minds blend them into seeing continuous light at 50Hz or above. Film projectors do this by flashing the light three times before moving onto the next frame [88]. At frame rates of 24 fps (typical movie projection frame rate), this technique means that light will flash 72 times per second, passing the flicker fusion threshold. Although inventing a mechanism to replicate this for a zoetrope would be difficult due to the spinning of a heavy wheel rather than a roll of film, this would alleviate the flickering problem and create a more accessible experience. As mentioned in Section 3.2, the praxinoscope uses mirrors rather than slits or a strobe light. This would get rid of the strobe problem but would reintroduce our original problem of looking at a 2D projection and not allow for interaction with light.

Another way to allow for more storytelling in a periodic zoetrope is by taking advantage of the fact that the wheel is not spinning that fast. The viewer sees each frame for a split second and the wheel rotates. The viewer does not see that specific frame for another 1-2 seconds (depending on the frame rate), so potentially, set pieces or backgrounds could be changed within that 1-2 seconds. Thus, by the next time the viewer sees the frame, the story is somehow different. This could be used to change the setting of the story or to have an object or character appear to teleport to different locations in the scene or even have objects disappear, to name a few examples.

Incorporating different types of light, such as UV or coloured lighting, could further enhance the storytelling capabilities of interactive zoetropes. UV light could reveal hidden elements, patterns, or details that are invisible under normal lighting, adding a layer of mystery or surprise to the narrative. Coloured lighting could be used to represent different emotional tones, time periods, or shifts in the storyline, allowing viewers to experience the same scene in new ways depending on the colour of light being projected. Additionally, borrowing the techniques used in 3D stereoscopic projection, like wearing glasses with polarizers, could let the animator encode more information into the scene. For example, maybe the viewer starts by wearing the glasses, but when they take them off, a different element of the story is revealed.

Ultimately, there are various creative ways to pack more storytelling into a zoetrope, and we hope to see this art form evolve further.

CHAPTER 4

Animatronics

4.1 Introduction: Democratizing Animatronics for Educational Storytelling

> Making is the most powerful way that we solve problems, express ideas and shape our world.
>
> Daniel Charny, Power of Making

Stories are the medium by which we decode the human experience. Storytelling has been discussed in the context of education as a way of cultivating imagination, empathy, and reflection of the world [25]. Through writing and performing, students build a literacy of storytelling which lays the foundation for deeper, more complex engagement with the world.

Animatronics is the art and science of building physical robotic puppets to bring a story to life with sound and motion. Animatronic shows have become a fun and popular attraction in theme parks, restaurants, and museums. This expressive practice is a unique combination of creativity and STEM (science, technology, engineering and math), which makes it a useful tool to engage students in both storytelling *and* robotics.

However, this expressivity comes with a high barrier to entry, and the medium is typically accessible only to trained engineers with ample resources. To help lower this barrier and make animatronics accessible to a wide range of ages, abilities, and socioeconomic status, we introduce an affordable yet versatile Paper Animatronics Kit for K-12 students to create papercraft puppet shows. The design of our kit is informed by the "critical making" movement established by thinkers such as Ratto and Garnet to describe the process of creating artifacts to explore and understand social and cultural issues, blending engineering, design, art, and social sciences [55], as well as Resnick's idea of "tinkerability," which he defines as "a playful, exploratory, iterative style of engaging with a problem or project"[119]. Critical making and tinkerability are not focused on the product of the making but the process of getting there which acts as a way for the maker to explore the world. By centring storytelling and the creative process, and allowing open-ended exploration with the technology, we aim to empower student agency and voice. Our kit thus uses an inquiry-based approach that taps into students' genuine curiosity about the world to solve problems, scaffolding the technical elements while allowing them to tinker and inject their creativity and identity.

In this chapter we describe the components of our kit in more detail, present the findings of our pilot studies, and end with key implications of using our animatronics kit in K-12 classrooms.

4.2 History: Evolution of Mechanical Puppetry

Although animatronics seem to be a part of modern technology, the idea of mimicking natural movements has been around since the ancient Egyptians. Priests would make large sculptures of the gods, sit in the shoulder joint, and move the arms. By those watching, this movement was perceived to be the gods embodying the sculptures. Later, in the 16th century Hans Bullman created the first androids that played musical instruments, Johann Muller created an artificial eagle, and John Dee created an articulated wooden beetle that could fly.

Although the previously mentioned models were very impressive for the time, in the 18th century, inventors were able to advance the world of moving sculptures when the use of clockwork became popularised. The gears, springs, and weights used inside this style of time piece allowed inventors to create models that could mimic human and animal movements. In 1737 Jacques de Vaucanson debuted the Transverse Flute Player [108], a humanoid figure that was able to play the flute. With the use of clockwork he was able to make the figure exhale and move its fingers. It was also able to play sound with the use of a carved drum that controlled the sound and volume. His second invention was a duck that was able to eat, drink, paddle, and digest to show the audience how ducks functioned. In 1770 Henri-Louis Jaquet-Droz and company created the most realistic models to date. The first was a child that could write and draw while their eyes followed the pen on the paper. The second was a pianist could move his fingers between keys and move his gaze from his fingers up to the music [108].

Figure 4.1: SAM Animatronic by Ihnatowicz.

As time went on, the models that inventors and artists created became more detailed and more realistic. It wasn't until the 20th century that the term *robot* was coined and the computer was invented, that they were able to create interactive and lifelike models. Figure 4.1 shows one of the first programmable cybernetic sculptures: SAM – Sound-Activated Mobile – created by Edward Ihnatowicz in 1968[153].

SAM is Ihnatowicz's first attempt at creating an articulated sculpture controlled by an electrical system. It is a flower-shaped sculpture made up of an aluminum stem and fiberglass sound reflector. Inside the stem there are small hydraulic pistons that allow the individual vertebrae to twist horizontally and vertically. On the fiberglass sound reflector there are microphones arranged in pairs and spaced so the computer can determine the location of the sound to allow for the sculpture to follow the audience as they walked around it. As an evolution to SAM, Ihnatowicz started work on his next project, The Senster, debuting in 1970 and is considered to be the world's first computer-controlled sculpture that would react to the audience. It has been described as a giant lobster claw with microphones instead of pincers, standing at 2.4 meters tall and 5 meters long. It is constructed out of welded steel tubes, hydraulic rams, actuators and an array of microphones. There are six joints along the arm with hydraulic rams and two custom actuators in the head controlled by a Phillips P9201 computer allowing for quicker sophisticated movements [153]. The Senster was able to move its head to home onto where sound was coming from, and the body would appear to follow instinctively. With the addition of a doppler radar the sculpture was able to react to the movement of the audience. For example if an audience member moved quickly The Senster would withdraw as if it were frightened.

Around the same time George Devol patented the first programmable robot that used a "magnetic process recorder" to determine how the robot moved [30]. This would lead to advancements across the board. This allowed for the models to do multiple actions at the same time, rather than having to do them in a sequence. This also made it possible for people like Walt Disney to create attractions that have been captivating audiences for generations. In 1960 Disney coined the phrase "audio-animatronics"; it refers to the use of sounds as cues for the next action [137]. Disney's first attempt at creating an interactive figure was President Abraham Lincoln for a future attraction. Although the company did not find it interesting, Robert Moses, the head of the 1964-1965 World's Fair, wanted it for the exhibition [46]. The only issues were that it wasn't finished and had some kinks and Disney was not sure how the audience would react. To figure them out, he created the attraction the Enchanted Tiki Room, added to Disneyland in 1963. The Enchanted Tiki Room consisted of 200 singing robotic birds, flowers, and tiki gods; the lifelike movement and sound captivated the public, giving Disney the confidence to continue working on President Lincoln and expand even further.

Figure 4.2: The Audio-Animatronic Great Moments with Mr. Lincoln displayed at Disney's Hollywood Studios.

During the 1964-1965 New York World's Fair, Disney had four attractions for different sponsors that helped him explore what audio-animatronics could do in animal and human form. The first was President Abraham Lincoln for the state of Illinois, the inside of which is shown in Figure 4.2. The Ford Magic Skyway involved the audience riding through the age of the dinosaurs where they could peek into the life of a dinosaur or caveman. The third is an attraction still being expanded on today, "It's A Small World – A Tribute to UNICEF," where viewers travel the globe singing with the children of the world [142]. The final attraction, General Electric's Carousel of Progress, shows the evolution of technology. With these first successful applications, Disney saw how useful audio-animatronics could be used not only to tell a story but to engage the audience as well.

So much progress in complexity of animatronics has been made, but because of this, creating animatronic characters is a complicated and technically advanced task. Such an engaging and expressive art form should be more accessible to people without all the resources of a big company.

4.3 Related Work: Animatronics & Education

We situate ourselves within a wider ecology of researchers, educators, artists, and makers working at the intersection of creativity, technology, and pedagogy.

4.3.1 STE(A)M, and the Creativity Gap

Of central importance to us is the so-called "creativity gap," defined as "an incongruity between the ostensible value educators place on creativity and its absence in schools" [125]. The creativity (or creative participation) gap manifests in education in a variety of ways: the removal of creativity from "academic" subjects and its partitioning into separate arts programs, the chronic underfunding of said programs [36], but also the inequitable access to the experience, skills, and tools required to flourish creatively in an era of digital media [68].

Emerging as an augmentation to the interdisciplinary field of STEM, STEAM aims to address the creativity gap by integrating the Arts into STEM education, emphasizing creative and design thinking as well as problem solving [73]. But this approach has been problematized by some, such as Mejias et al. [93], who argued that STEAM education fails when either art or engineering takes precedence over the other; introducing STEAM in a nuanced way proves challenging. Liao, on the other hand, argues for a transdisciplinary, arts-integrated approach which is centered on "the creation of art that is simultaneously applied work" [84]. Motivating STEM tasks with creativity promotes student voice and choice, which inspires students to make something they truly care about and are proud of.

4.3.2 STEAM Education in Action

The implementation of STEAM education varies; some programs aim to integrate robotics into the classroom, which requires planning and teacher training. The "Arts & Bots" program, for instance, aims to increase empowerment and inclusion in STEM disciplines by integrating robotics into middle-school core subject classes [82], pairing the Hummingbird robotics kit with a custom software programming interface. The Arts & Bots team has conducted a series of user studies with teachers and students to uncover student learning outcomes and attitudes toward STEM [31] and the challenges teachers face when planning and implementing the program [50]. Although teachers were mainly successful in their attempts to combine robotics with their course material, they found that the nature of subject-specific curricula constrains teachers' pedagogical choices, limiting opportunities for open-ended storytelling. Additionally, controlling the robots requires coding skills, adding to the teacher training required [51].

Many commercial robotics education kits exist, such as Hummingbird [101], used in the Arts & Bots program, which targets Grades 4-12 and includes compatibility with the micro:bit, and littleBits [12], a system of modular pre-assembled circuit boards which snap together magnetically, making it accessible for young children. These kits usually include various sensors, servos, and LEDs. For example, Phidgets allow for more complex coding projects by providing modular building blocks of electrical components [63]. Although versatile, they are very expensive, generally costing over $1000 USD to outfit a classroom, making it infeasible for schools without the means to buy them.

Additionally, numerous after-school programs and summer workshops incorporate robotics into creative tasks, such as an upper elementary robotics program that designed an animatronic zoo [100] and a program for middle-school girls to build expressive robots [50]. However, while extra-curricular programs such as these motivate young people to engage in creative and design thinking, they are prone to the same inequities of access that drive the creativity gap [128].

4.3.3 Motivating STEM Through Animatronics

While much research has focused on integrating creativity into STEM through visual art, another compelling avenue for student expression is through stories. Incorporating puppets into storytelling is a way of "making the story come alive" and can give kids new perspectives on, and relationships with, stories [21]. Animatronics and puppetry provide a natural entry point for the integration of STEM with storytelling and creative thinking. Robotic puppetry has been found to engage children in storytelling by allowing them to take an active role in the story, for example, by inviting kids to interact with the puppets during the story [74] or having the kids do the puppeteering themselves [83]. Huang et al. described a 5-day workshop with 11-13 year-olds who created interactive "e-crafts" and accompanying written stories [59], finding that storytelling allowed students to inject their own interests and identities into their learning, thereby deepening their engagement with the STEM and design tasks. Alford et al. used robotics to combine STEM with drama, hosting a 3-day animatronics workshop where high school drama students wrote plays and built and programmed their own robot actors [2]. Their workshop, while very demanding in terms of the materials and expertise required, demonstrated the potential for animatronics to serve as an outlet for children's creativity.

There is a need for affordable technical supplies in order to participate in animatronics or any of the mentioned STEM tasks. Papercraft has been explored as a low-cost medium for students to experiment with robotics. Systems like AutoGami [152] and FoldMecha [104] provide software for students to design and program moving paper creations. While AutoGami has been used for more representational papercraft, artifacts made with FoldMecha are more like paper automata with electronic actuators. MoveableMaker software allows for easy creation of interactive papercraft, where a user moves one or more elements to generate an effect [4], while MakerWear takes a hardware approach and provides students with a tangible modular construction kit with which to make creative computational wearable artifacts [72]. These kits challenge kids to come up with creations involving electronics, but none focus on storytelling, allowing the creative elements to take a back-seat to the mechanical challenge.

4.3.4 Collaborative Making

Dieter and Lovnik, in their "Theses on Making in the Digital Age," state that "[t]he maker is always plural. We all know we never make things alone... We feel a constant pressure to invent and discover new tools to support collaboration" [33]. Cross-age peer mentoring is a collaborative model which has been explored in education as a means of mobilizing student knowledge and building social-emotional skills. Students also benefit from developing friendships, gaining confidence, and strengthening knowledge and skills [11].

Boling et al. explored cross-grade mentorship in outdoor education, pairing students from Grades 6 and 3 during a field trip to study water quality of a local river. They found that engaging in mentorship "deepens interest, investment, and ultimately ownership of new learning" for both the mentor and the mentee [71]. In the context of STEAM education, Tenhovirta et al. studied cross-age

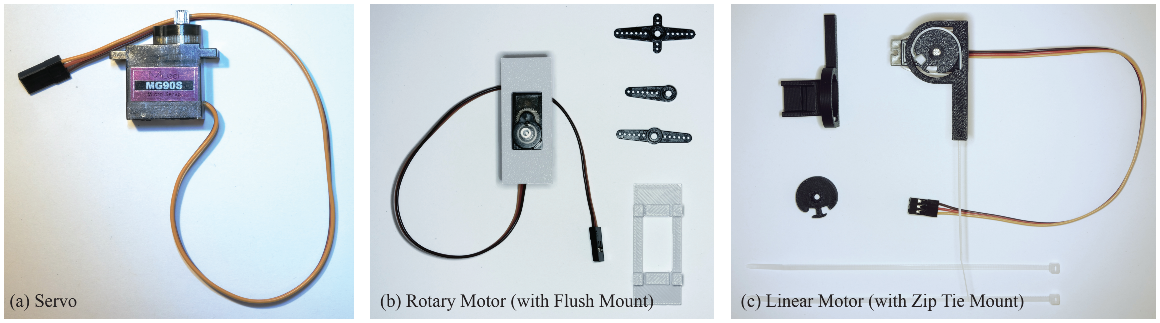

Figure 4.3: The servo (a) is about 30 x 30 x 12 mm. The Flush Mount (b) allows for rotary motion. The Zip Tie Mount (c) allows for linear motion.

tutoring in a maker-centered lower secondary school, examining mentor/mentee relationships within teams of students working on a collaborative design task [132]. They found that "young people have impressive sociodigital skills that could provide valuable social learning resources when their use is legitimised through peer tutoring practices."

4.4 An Accessible Animatronics Kit for Students and Teachers

One of the key contributions of this work is the Animatronics Kit, a low-cost educational kit combining creativity and STEM tailored for use in elementary classrooms. Our kit makes it easy for kids to create talking paper robots of their own design where motions are synchronized with sound. These robots do narrative storytelling, and their mouths move convincingly as they speak. We analyse the kit's effectiveness from a high-level user experience perspective, i.e. we explain how the parts work – their input and output, show what students can build with them, and examine the usability of the parts and completeness of the kit.

Because it is specialised for the assembly of an animatronic show rather than providing more generalised functionality, this focus allows the kit to be more affordable and streamlined compared to mainstream robotics kits. Along with ease of use, cost is a major factor in accessibility. Our most expensive board costs about 7 US dollars to produce in 50 unit quantities, allowing us to offer these components at prices that are highly competitive with inexpensive Arduino-based kits that have found wide adoption in schools.

Hobby servos are commonly used in model airplanes to move control surfaces. They are complete servo mechanisms incorporating a motor, a gear train to trade off speed for torque, a position sensor to measure the current angle of the output shaft, and a small printed circuit board which receives position commands and provides power to the motor to move towards the commanded angle. Micro servos, such as the Hobby King HK1615178, Miuzei MG90S, or Feetech FS90, are inexpensive (typically less than 5 US dollars each) and have sufficient speed and torque for actuating paper mechanisms that mimic mouth movements. We use the Miuzei MG90S 9G Micro Servo model for its affordability, small size, and reliability, as seen in Figure 4.3 (a).

Each kit consists of one Linear Motor, one Rotational Motor, three printed circuit boards (PCBs), and a battery pack. We also provide template animal characters from Woo! Jr [120] printed on 8.5 x 11" heavy-duty paper, as well as double-sided tape to attach the Motors to the cardboard or

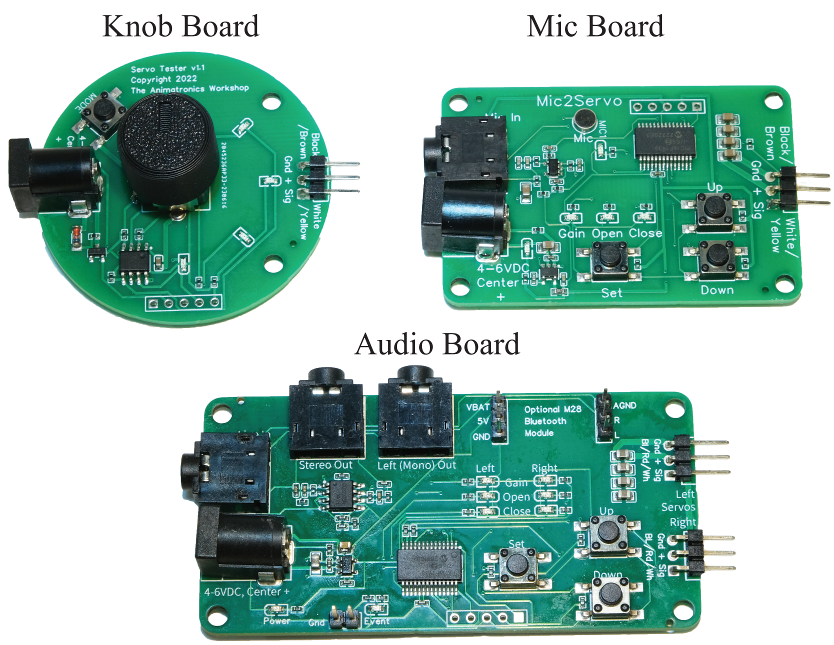

Figure 4.4: Our boards allow users to control the motor in different ways.

paper puppets. We assume classrooms have access to common crafting materials and tools, such as scissors, glue, card stock, etc.

Each of our Motor units comes pre-assembled and comprises a small electronic servo, a type of motor common in hobby applications, such as robotics and model aircraft, and a custom-designed 3D-printed mount to allow easy attachment of the servo to the puppet (see Figure 4.3). The Rotational Motor unit includes three different horns (in black) that clip onto the shaft, which provide a small surface area to attach the moving element of the puppet. This allows the user to create swinging motions, such as a waving arm or a kicking leg. The Linear Motor mount includes a small mechanism to convert the servo's rotational motion into linear motion, with a zip-tie that winds and unwinds to attach the moving element. This motor is intended to allow the creation of talking characters whose mouths move up and down, but it has many more applications, such as as a punching arm, or even a grasping claw.

To control the motor, we have designed three circuit boards (see Figure 4.4) customized specifically for the purpose of animatronic storytelling through puppetry. They are intended to scaffold students' understanding of the features and slowly introduce new capabilities.

The three boards work similarly in terms of connecting and configuring: the motor plugs into a 3-pin connector on the right side of each board, and a battery pack with four AA batteries plugs into the jack on the left side of each board, as seen in Figure 4.8. Buttons on the board allow the user to adjust the range of motion by setting the maximum "open" and "close" positions, which are saved for subsequent use, even after the boards are turned off and back on.

The boards differ in how they allow the user to control the motion of their puppet. The first and simplest is the Knob Board, which lets the user directly control the motor shaft angle using the knob. It also has a "sweep" mode, which automatically spins the motor back and forth between its full range of motion with variable frequency.

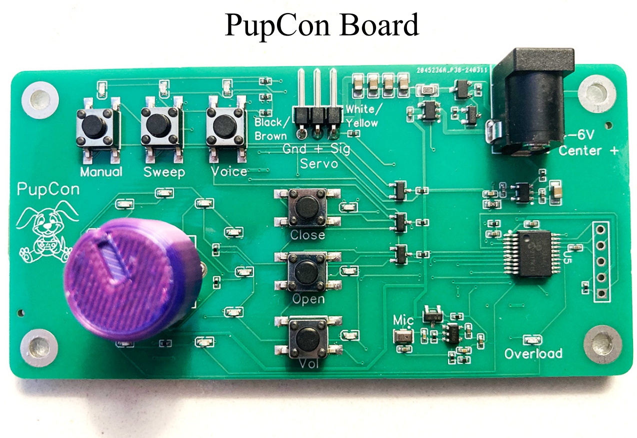

Figure 4.5: The PupCon Board (puppet controller) board combines the functionality of the Knob Board and the Mic Board.

The Mic Board uses an onboard microphone that the user can talk into, rather than controlling the motor using a knob. The motor moves proportionally to the volume of the audio input, allowing the user to perform live shows with the puppet. This board additionally allows the user to configure the gain, the amount the motor moves for a given volume of input, which can help when the puppeteer is in a noisy classroom.

After our pilot study with the Junior Kindergarten students, described in Section 4.7.1, we observed that young students like to press all the buttons on the board at the same time. We also discovered an expected confusion about the difference in interface for the Knob Board and the Mic Board. The Knob Boards uses the knob itself to not only control the motor but also to change the motor limit settings, whereas the Mic Board uses the microphone to control the motor but uses Up and Down push buttons where the user has to press and hold the button to change the motor limits and microphone gain settings. The user can visually see the motor limit changing while they hold the button, but there is no visual indication on how much the microphone sensitivity changes when they hold the button. The lack of consistency between the boards, along with the redundancy in functionality led us to design the PupCon Board, as seen in Figure 4.5 that combines the boards and unifies the user interface by using the knob for changing settings *and* the motor when it is in Manual Mode. In Voice Mode, the PupCon Board works the same way as the Mic Board and uses a higher quality microphone. We do not use these combined boards in our user study but will be used in future studies and activities to replace the Knob and Mic Boards.

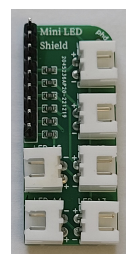

Figure 4.6: The LED shield.

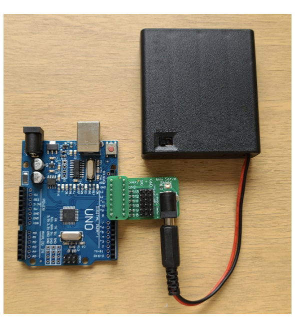

Figure 4.7: The servo shield.

Most complex is the Audio Board. Similar to the Mic Board, this board responds to sound but includes a 3.5 mm audio input jack in place of a microphone, allowing students to play recorded audio to create pre-recorded scenes and skits. The board also includes an audio output, allowing the user to pass the audio to an external speaker when presenting. This board additionally includes two 3-pin connectors, allowing two motor units to be controlled simultaneously and independently. For this study, this feature was simplified to ensure accessibility for young children by making the two motors move in unison.

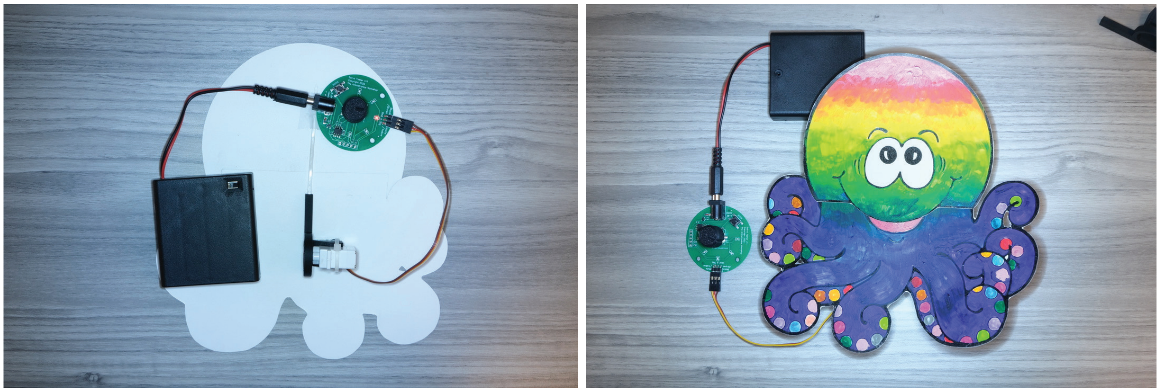

Figure 4.8: An assembled puppet.

The Audio Board also supports extension to an Arduino microcontroller, incorporating programming tasks into the show. The user can record dialogue on the right audio track and use the left audio track to control cues during the show, triggering events to happen with a series of high signals, or "beeps." Extending the kit further, Figures 4.7 and 4.6 show two more boards that conveniently attach to an Arduino. The servo shield (Figure 4.7) plugs into the pins on one side of an Arduino Uno and provides 6 more slots to plug in more servos. This can be used in conjunction with the Audio Board, or without. To power the servos a battery pack can be plugged into the servo shield. This is a separate power source from the one powering the Arduino itself *and* the one powering the Audio Board if it's also being used in the project. The LED shield in Figure 4.6 works similarly to the servo shield, and plugs into the pins on the other side of an Arduino Uno (so both shields can be used at the same time). It allows up to 6 LEDs to be plugged in and programmed with the Arduino. The shields are meant to provide even more customisation and complexity to an animatronics show using popular microcontrollers, making it accessible but also adding challenge for students.

4.5 Research Goals: Animatronics in the Classroom

In this chapter we categorize our research goals in four main areas:

1. Exploring the challenges faced during an animatronics task for varying age groups
2. Determining how our kit design choices impact student creativity
3. Exploring the impact of mentorship on task independence for younger students and understanding for older students
4. Assessing the feasibility of having non-tech teachers lead technical activities, specifically an animatronics assignment

Table 4.1 shows the four workshops we will discuss in this thesis, all of which took place over the last year. Two Robotics Camp Workshops bookend our School Study.

Age and Challenge Level Across our studies, we worked with students ranging from 4 to 16 years old, each with varying skill levels and prior experience. Our animatronics kit is designed to accommodate this wide range of abilities, offering tasks that scale from simple puppet construction to more advanced activities such as using the Audio Boards and integrating Arduino programming to control the puppets and make more complex stories. We anticipated that while younger students might engage primarily with the basic mechanical aspects, older students could potentially explore the more advanced features of the kit. A key goal of these studies was to observe how far students could progress, based on their age, skills, and the challenge level of the tasks presented.

Kit Design Choices The animatronics kit was designed to be both accessible and adaptable, with components that allow for a wide range of creative possibilities. Our primary objective was to provide students with intuitive parts that could be easily assembled to create functional puppets, while also giving them the flexibility to experiment and personalize their designs.

We evaluated how well the kit's parts functioned across various projects. We wanted to understand where the kit's design facilitated smooth and intuitive building, and where it presented challenges that limited what students could achieve. These findings – whether related to part durability, ease of assembly, or adaptability to different creative ideas – will inform future revisions of the kit to better meet the needs of young makers.

Impact of Mentorship Storytelling is both individual and communal, so given the wide age range of students, we explored mentorship as a way to support learning. By pairing older or more advanced students with younger ones, we aimed to see if this dynamic would help younger students become more independent or successful, while also reinforcing the older students' understanding through teaching. Our goal was to assess whether this approach improved outcomes for both groups and how it might inform the design of future workshops.

Non-technical Teacher Experience While technical educators, such as those leading our robotics camps and the school's technology teacher in the JK pilot study, quickly adapted to using the animatronics kit, we wanted to explore whether non-technical teachers could similarly engage with the kit in their classrooms. This was a critical area of study, as it would determine the kit's broader applicability across diverse educational settings.

We wanted to explore how non-technical teachers perceived the kit's usability and what aspects, if any, caused them concern. Additionally, we were interested in understanding what types of support might help make the process smoother for them. By examining their experiences, we hoped to gather insights into how the kit can be made more accessible and adaptable, ensuring that teachers without a technical background can confidently lead animatronics projects.

In the following sections we describe each workshop and our findings and discuss results in the context of our research goals.

4.6 Robotics Camp Workshop 1

A private robotics institute in the greater Toronto area that offers after school activities and summer camps ran a 1-day animatronics workshop using our kit. The workshop was hosted at the University of Toronto. Thirty-two students from the ages of 8-14 years old, who were already participating in a week-long robotics camp, came to the 6-hour workshop, which was split into three hours before

Study	Robotics Camp 1	JK Puppets	Grade 2 & 6	Robotics Camp 2
Age level (yrs)	8-16	4-5	6-7 & 11-12	7-16
# of students	32	30, 4 at a time	22 & 24	34
Length of study	1 day	3 weeks	8 weeks	1 day
Boards used	Knob, Mic	Knob, Mic	Knob, Mic, Audio	PupCon, Audio
Motors used	Linear Motors	Linear Motors	Linear & Rotary Motors	Linear Motors
Other parts used				Servo Shield
Activity	Construct puppets from templates, puppets from photos	Sketch and construct original character puppets, write short story	Learn boards, write stories, construct original puppets, perform live animatronic shows	Construct puppets from templates, programming task with Audio boards and Arduino
Age & Challenge takeaways	Large age range is challenging for curriculum planning	Students as young as junior kindergarten can make puppets	Mentoring and inquiry-based learning are valid and successful ways of teaching animatronics	One day is too short to meaningfully incorporate storytelling into high level technical task
Kit takeaways	Sweep function is fun	Boards are confusing and redundant, kids like pressing all buttons	Mic board gain setting either not used or not sufficient	Scaffolding so students can work independently improves classroom management, programming success
Mentorship takeaways	No mentorship leads to frustration		Younger students were able to work independently after mentorship	Small groups of peer mentorship works well
Teacher takeaways	Run by tech teacher	Run by tech teacher	Non-tech K-6 teachers were open to trying, went well but could go more smoothly	Run by tech teacher
Curriculum takeaways	Would have liked to incorporate more creativity	Fine motor skills development and story writing are difficult but do-able with teacher guidance	Integrating animatronics into subject-specific curriculum requires planning and willing teachers	It is possible to begin programming tasks with robotics students in 1 day, but creativity and storytelling does not fit

Table 4.1: Over the course of this thesis, we ran four workshops with differing timelines, demographics, and goals. The table shows each of our workshops in chronological order from left to right, summarizing details about the age levels of the participants, which technology from our kit we used, what activity the students completed, and takeaways from different perspectives.

and after lunch. Four staff members of the robotics institute were there to manage the classroom as the students worked.

In the morning the research team gave a short 10 minute lesson about the boards and motor units in the kit. The lesson also covered how to make a puppet from a template by cutting it into two parts and attaching the linear motor to one side and the zip tie coming out of the motor mount to the other, as seen in Figure 4.8. Although templates were provided by the research team and available to students, we encouraged them to create original characters instead. Most chose to use a template, but four decided to make their own, which happened to be mostly characters from video games. Some students didn't want to make a character, so one instead chose to make his "puppet" a drawing of a gun from a video game, which showed his creative thinking and correct understanding of the motor's movement.

After lunch, students moved on to make puppets out of pre-taken photos of themselves, one with their mouth closed and one with their mouth open. The challenge in this task came from making puppets with three different parts instead of two. The students had to recreate the example in Figure 4.9. Making this puppet involves cutting the mouth and chin piece out of the picture of the student with a closed mouth. Students too young to use xacto knives had to wait for a staff member to do it for them, causing a back-up of progress at this step. Instead of paper, this puppet is also backed with cardboard for sturdiness, which is a tougher material to work with. The students then have to cut a hole in the cardboard backing and tape their chin piece to a smaller piece of cardboard. The linear motor attaches to the larger cardboard and the zip tie to the smaller piece, which is suspended between the two photos of the student.

This increase in difficulty in this session caused frustration with some students struggling to cut a hole in the paper and others not understanding the mechanics of the example puppet. Other students did not want to take their puppets from the morning apart to re-use the motor. Additionally, each student only had one print-out of their two photos, which meant that messing it up required some creative solutions, as there were no back-up photos to use.

When the workshop was planned, we expected 10-15 kids from the ages of 8-11, but the activity gained more popularity than we thought, and more parents signed their children up to participate. Because of the large age range, doing the same activity with all of the students did not work well. Students on the low end of the age range struggled to understand the boards and cut pieces necessary to make the puppets. Students on the high end finished early and became bored. The afternoon task took more time than expected and was more complex for the given age range than we realized.

However, the morning task seemed to be enjoyable and every student managed to complete it successfully. It also showed us that students really liked the "sweep" mode of the board, where the motor automatically moves back and forth its whole range of motion. This was useful information, as we were considering getting rid of it in the next iteration of the board. The troubles we faced in the afternoon session led us to emphasise mentoring in future studies in order to combat the older students finishing early and the younger ones struggling on their own, improving classroom management issues.

Figure 4.9: The assignment to make a puppet from photos of themselves requires cutting the bottom lip and chin out of the photo with a closed mouth and placing it on top of the photo with an open mouth. Here we use the Knob board to show the closed and open mouth positions of the puppet.

4.7 School Workshop Study

The goal of our school workshop study was to evaluate the kit's effectiveness in three key areas: the ease and comfort for teachers to implement it in their classrooms, the ease of use and expressivity for students, and the identification of mechanical or technical aspects that required improvement. We also aimed to explore the impact of peer mentorship and teamwork within the context of animatronics. We conducted a series of in-classroom user studies at an urban independent K-6 school affiliated with the authors' institution and located in a major Canadian city (*the lab school*). We had approval of our ethics protocol through the REB at our institution, as well as the board of the lab school in which we conducted the study. To preserve participants' anonymity, all names have been changed.

4.7.1 JK Pilot Study

Figure 4.10: A JK student's sketch vs. the puppet.

Working with the school's Tech Teacher, Ricardo, we began with a pilot study with a class of Junior Kindergarten (JK) students (aged 4-5) to validate the parts of the prototype kits. We established a collaborative relationship with Ricardo, who worked closely with the research team throughout our time at the lab school. Ricardo, who is also the school's Special Education teacher, has a personal interest in technology and helps integrate technology into the classroom on a case by case basis with all grade levels. The research team observed each stage of the pilot and only got involved during the part of the process where the students used boards.

He led the JK students through a series of lessons, starting from "What is a puppet?" and leading all the way to each student creating their own paper animatronic character. The construction part

Figure 4.11: JK student character sketches and resulting puppets before the students had to cut them and attach motors.

of the pilot took the form of three 1-hour sessions over the course of three weeks. In groups of three or four, the 30 JK students in the class worked with the research team to create paper puppets from sketches the students had drawn with their teacher. Figure 4.11 shows some of the JK students' puppets sitting on top of the reference sketches. During the creation process with students, some design details about the characters needed to change for mechanical reasons or because of the materials available (see Figure 4.10.

We assisted the students in cutting their puppet into two pieces so that we could help them attach the Linear Motors to one side and the zip tie to the other. Changing the settings on the Knob Board and Mic Board proved difficult for the 4-year olds, who had little patience and undeveloped motor skills, so Ricardo and the research team had to set the motor limits for them. The students then performed short stories with one line of dialogue using the Mic Board. They had no problem using both boards for their intended purposes once the settings were changed and saved. They even enjoyed watching their character talk when they spoke into the Mic Board. This pilot demonstrated to the research team that the parts in the kit not only were understandable for very young children but were also robust to rough handling. It was also fascinating to watch the students learn how to embody their character and grow their literacy and storytelling skills.

4.7.2 Grade 2 & 6 Study

Participants Our participants comprise the students and teachers of two classes at the lab school: a Grade 6 class and their teacher, Anita, and a Grade 2 class and their teacher, Sonny. The teachers signed consent forms to be interviewed before and after the process of facilitating the workshop. Parents of the students in the classes signed consent forms giving permission for data collection of audio, video, and photo of the students and their creations, and were given the option to blur photos and alter audio recordings of their child in publications.

Grade 6 teacher Anita had a background in Kinesiology and 22 years of experience teaching across grades K-8. Working initially as a physical education teacher before transitioning to being a classroom teacher, she completed professional development in reading and elementary science. In our pre-interview with Anita, she discussed her enthusiasm for cross-curricular integration in her teaching. Her class consisted of 24 students between the ages of 11 and 12 (12 girls, 11 boys, 1 non-binary).

The Grade 2 teacher, Sonny, had an undergraduate degree in history with a minor in biochemistry. After working as an outdoor education facilitator, he decided to get his Master's of Arts in Child Study and Education. He had four years of classroom experience, and also had worked as a physical education teacher, as well as an occasional teacher, before becoming a Grade 2 teacher in the lab school. He discussed with us the importance of understanding each of the unique learners in his classroom, and his teaching emphasized student voice and choice. The Grade 2 class had 22 students between the ages of 7 and 8 (11 girls, 11 boys).

The two teachers participated in a 45-minute unrecorded Teacher Training session in which they had the opportunity to explore the animatronics kit themselves. The research team went through each board, explaining how to access and change the settings and how to plug them into the battery pack and motor. The teachers each made an animated puppet character using pre-made characters printed on card stock. They were able to quickly grasp the idea of how to use the components in the kit to make a puppet. The research team left the kits with the teachers so they could continue tinkering with their characters.

The teachers consulted with the research team but were given a large degree of independence to introduce the animatronics kits into their classrooms in the manner they thought best. As the two classes had an established cross-grade mentoring system, each Grade 6 student being paired with a Grade 2 "special friend," the teachers were enthusiastic to incorporate this mentorship aspect into the workshop and collaborated closely when planning their instruction. This also informed our research question about mentorship and prompted us to craft interview questions for students and teachers to investigate its role in the experience.

Interviews and Group Discussions We conducted individual semi-structured interviews with both the teachers and the students. In order to get a sense of the background and interests of our teacher participants, we conducted 1-hour interviews (*pre-interviews*) with each teacher before beginning the study. At the conclusion of the workshop, once most of their students had completed and presented their puppet shows, we conducted 30 minute interviews (*post-interviews*) with the teachers to debrief them and hear their observations and feedback on the workshop and the kits, as well as their suggestions for improvements.

In addition, we collected feedback from the students through two 10-minute, group discussions with the Grade 6 students and a number of short, individual or small-group semi-structured interviews with a handful of students from each class, each lasting approximately 5-10 minutes. Students were selected to be interviewed from those who had finished creating their puppet shows on the basis of teacher recommendations and student interest (see Section 4.14 for the list of interview questions we used for teachers and students throughout the study). In total, we spoke to nine Grade 6 students over six interviews, of which two interviews comprising four students were discarded, and ten Grade

2 students over seven interviews, of which none were discarded.

Data Collection Throughout the School Study Workshop, we collected a variety of multi-modal data. During all sessions, students' work was documented through photos and videos showing their creative process and work in progress. Researchers circulated among the students to observe, provide support, and chat informally with them about their process and their experience with the kit. In addition, we collected video and audio of the whole room during Grade 6 group discussions, but the video portion of this data proved unusable and was discarded.

All data were anonymised and stored digitally on a secure cloud storage service, and only those members of the research team directly involved in the data processing and analysis were given access. Audio recordings were automatically transcribed securely on the researchers' device using OpenAI's Whisper algorithm [116]. These transcriptions were then manually verified by a member of the research team.

Coding and Thematic Analysis We performed an iterative process of coding and thematic analysis [95, 102] on the transcripts from the interviews and group discussions. Two researchers independently performed two rounds of open coding on the transcripts, each followed by discussions to ensure inter-coder agreement. These codes were then analysed over three collaborative sessions, in which the research team reviewed all the codes and artifacts to identify salient themes which we developed into the key implications discussed in Section 4.9.

4.8 School Workshop Findings

For our main study, we worked with two classes, one Grade 2 and one Grade 6, over 13 sessions across a period of approximately eight weeks. We organize this section by groups of sessions with each grade, including the two special friends sessions when they worked together.

Session 1 - Grade 6 Exploration In the first 1.5 hour session, Anita, supported by Ricardo, introduced the kit to her students using an inquiry-based approach. The Grade 6 students each were given a Linear Motor, a Mic Board, and a battery pack, and tasked with independently figuring out how to assemble the parts and exploring how the kit works, including using their voice to actuate the motor, and changing all three settings on the board: both motor limits and the microphone sensitivity. Once students were familiar with the basic functionality, they were then asked to create a character out of paper. Students were given the freedom to create whatever they wanted, and they found inspiration from many sources, basing their puppets on animals, characters from popular culture and even each other. Anita chose to offer students struggling for inspiration the option to use template characters provided in the kit, but only two students chose this. Anita did not provide this option going forward, preferring to encourage students to create their own character. Some found the suggestion to make a talking character limiting, and we saw many creative applications of the Linear Motor in papercraft mechanisms, such as Leo's Whack-a-Mole game or Ryan's face with an animated tongue (see Figure 4.13). In the busy classroom, some students grew frustrated at the lack of control offered by the Mic Board, which they found to be too sensitive for the noisy classroom environment. We thus offered all students the option to use the Knob Board, which provides more

Figure 4.12: Grade 6 student Sasha creates a puppet with her special friend.

direct, manual control. After students had built their puppets, they performed improvised skits for the class, either alone or in groups.

To close the session, Ricardo led a group discussion asking students what they found enjoyable and difficult. Students described their approaches to design and construction. One student, Francis, discussed the creative compromise he had to come to when assembling his puppet, saying "We wanted people to see the full drawing, so we ended up putting half of [the motor] actually showing." Another student, Sasha, reported frustration with the Mic Boards due to latency and sensitivity, recounting that "I would have to have people around me be quiet so that it would stop moving and it would still go from little sounds." Despite this, she enjoyed the ease with which she could bring her vision to life, continuing, "It was also really satisfying and easy that all you have to do is take the machine and just plug things in and make a puppet." Because of the inquiry-based approach, there was almost no direct instruction about the boards' features, and demonstrably many students did not figure out how to alleviate the noise problem by setting the microphone gain to be lower.

Session 2 - Special Friend Teaching In this 40-minute session facilitated by Anita and Sonny, the Grade 6 class was joined by the Grade 2s, and students broke off into *Special Friends* groups, pre-assigned cross-grade mentorship pairings. Each group was given a battery pack and could choose to use either the Knob Board or the Mic Board, and either the Linear Motor or the Rotational Motor. Students were tasked with inventing a character and turning them into an animatronic puppet, with the Grade 6 students guiding their Grade 2 partners on the use of the electronic components, as seen in Figure 4.12.

After 30 minutes, the Grade 2s returned to their classroom, and one of the researchers led a group discussion with the Grade 6s reflecting on mentorship experience. We asked the Grade 6s which of the boards they used with their special friend. Most students who responded said they ended up using the Knob Board ("the twisty one") over the Mic Board ("the talking one"). Consistent with their experience from the previous session, students found the Mic Board hard to control, with

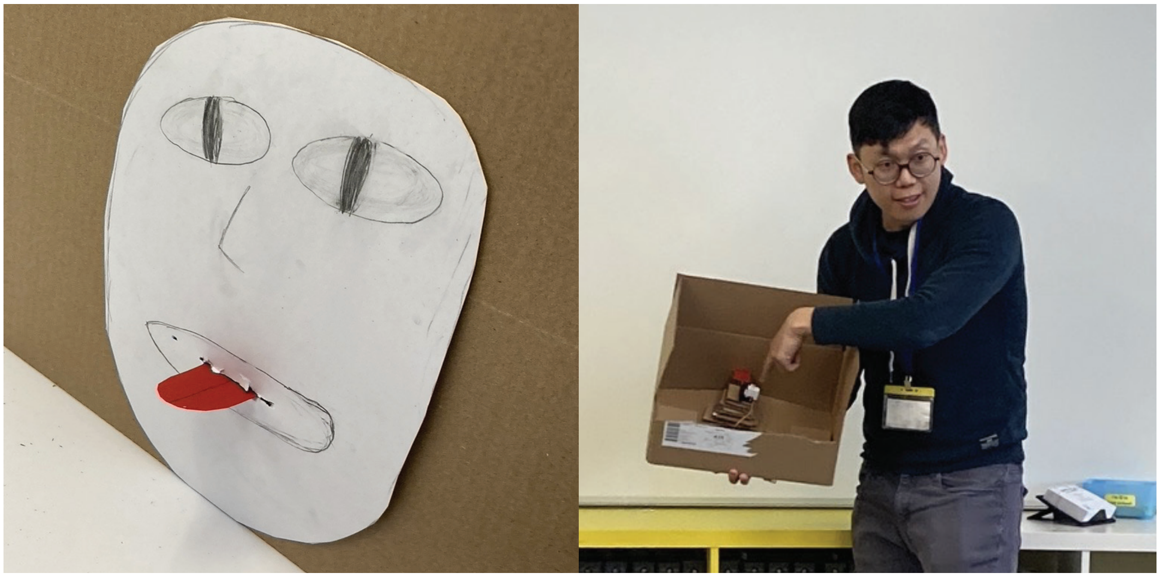

Figure 4.13: During a class discussion, tech teacher Ricardo highlighted the problem solving steps Grade 6 student Ryan followed to make the tongue of his puppet move in and out using the Linear Motor.

a student reporting, “It would respond too late and it would make random movements.” Another student, Cindy, recounted that she initially used the Mic Board but switched to the Knob Board because the extra buttons distracted her special friend. Students reported that the Knob Board gave them more precise control and that they could more easily understand the correspondence between their input and the resulting motion. In addition, some students preferred the Knob Boards because of the automated sweep feature. However, the Mic Board seemed to spark a particular sense of wonder in the younger Grade 2 students. Sasha picked it because her special friend wanted to try it:

> I think it was more fun for her, because she got to– I don’t know. She just really enjoyed getting to speak and seeing its mouth move, and I think it was just kind of cool. The twisty one, it makes a lot of sense, it’s like you twist this and it goes up and down. But the talking one is more magical and fun when you talk and it talks.

Sessions 3-5 - Grade 6 Writing and Puppets Over these three one-hour sessions, the Grade 6 students worked individually or in pairs to develop a final puppet, choosing either to further develop the puppet they created with their special friend or to develop a puppet based on a character from the stories they had been writing in their literacy class. The sessions were led by Anita, and Ricardo was present at one of these sessions to provide additional support.

As the students’ confidence with the kit grew, so did the sophistication of their creations. Working with characters in which they were already invested, students employed creative design to create puppets which fit into settings of their own invention. Leo described how he choose to bring the giant worm from his story to life: “I picked which character would look good moving, and I thought about which characters would be able to move easily.”

Figure 4.14: The progress of making a puppet, from sketch to show. Cindy shows her special friend the spaceship from her story.

Figure 4.15: Grade 6 student Hiro records audio to control his puppet.

Students who completed their puppet early and were seeking more challenge were invited to use the Audio Board to create an animated skit. Using a classroom laptop and an online voice recording service [133], students were able to record lines of dialogue which, when played back through the Audio Board, animated the puppet (see Figure 4.15). While experimenting with playing different audio through the board, Leo and Hiro tried playing music, hoping it would look like the puppet was singing but were disappointed to find that the puppet's mouth simply opened wide when the music played. Leo later recounted how this experience deepened his understanding of the technology: "With the song, it's not speaking the lyrics. It's just open when there's noise, . . . and then it's shut when there's no noise."

While some students, such as Leo and Hiro, put their efforts into developing the technological complexity of their puppets, other students focused on refining the artistic components. Cindy spent the majority of her efforts carefully drawing her cartoon rocketeer (see Figure 4.14). She used the Knob Board's sweep function to automatically animate the rocket's flames.

Figure 4.16: Simone presents her finished puppet show to her special friend.

Session 6 - Grade 6 Puppet Show and Tell In this 30-minute consolidatory session, which concluded the Grade 6 students' involvement in the workshop, the Grade 2 class once again joined the Grade 6s, two weeks after the initial mentoring session to see their special friends' finished puppet shows. The Grade 2 students had already begun brainstorming stories and characters for their own puppets at this point. This session gave them an opportunity to draw inspiration from their Grade 6 peers. At the same time, the Grade 6 students had the chance to show the culmination of the efforts in creating their puppets as they brought them to life for their special friends. The Grade 6s used their puppets, as well as voices and gestures, to breathe life into

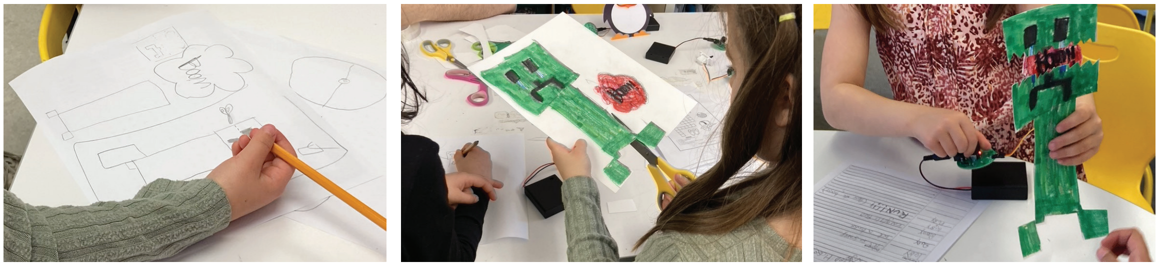

Figure 4.18: The progress of Grade 2 student River's Minecraft creeper puppet character named "Boomy McBoomerface."

their stories. Some students used an online digital storytelling tool, StoryJumper [130], to enhance their performance.

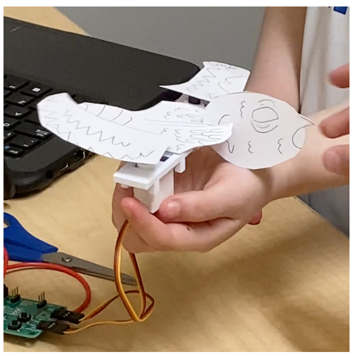

Figure 4.17: Ryan's bird.

Near the end of the session, Grade 6 student Ryan created an animatronic bird with a few others, at the suggestion of his special friend. Telling us that he "was used to microcontrollers," he made use of the Audio Board's two motor outputs to connect a pair of Rotational Motors to act as flapping wings. Other students contributed artwork for the bird's wings and head, and Ryan recorded an onomatopoeic sound effect, "*flippity-flappity*," to animate the wings (see Figure 4.17), delighting the Grade 2 students as they flapped back and forth.

Sessions 7-13 - Grade 2 Puppet Shows For the remainder of the workshop, the Grade 2 students used the kits independently, working on their own puppet shows. For each of our sessions with the Grade 2 students, facilitated by Sonny, we worked with half-groups of 11 students at a time. Ricardo was present at two of these sessions for support and observation. We worked with each half-group three times over six 1-hour sessions, supporting them as they developed their characters into puppets. Group instruction was largely the same between half-groups, and students worked at various paces in pairs or small groups.

Sonny began with each half-group by (re-)introducing the components of the kit, asking whether they remembered their names and how they connected together. Despite the fact that the Grade 2 students had not received direct instruction on the kit, they were able to describe the purpose of the motor, board, and battery, and how to connect them, due to the hands-on experience with their special friends. They were easily able to recall how to use the boards they had familiarity with (either the Knob Board or the Mic Board) and were able to immediately begin creating puppets. They were tasked with picking a character from their story and bringing it to life.

Compared to the Grade 6 students, who generally only needed help when facing technical problems, such as flat batteries, the Grade 2s needed more teacher support. In particular, attaching the motor to the paper cutouts proved taxing to the students' fine motor skills. This was partly exacerbated by the small size some students drew their characters on the card stock. After the first pair of sessions, Sonny chose to make an exemplar puppet available to students based on one of the templates provided with the kit to give them an idea of sizing.

As the students began completing their puppets, Sonny had them form small groups of 2-4 and

Figure 4.20: Rose and Bernadotte write a script with lines for their troll character.

collaborate to write a cross-over script where each of their characters meet. As we saw with the Grade 6 class, students engaged with the kit in different ways; some were driven by the making, such as Amir, who was determined to build a tripod to allow his alien puppet to stand while it talked, as seen in Figure 4.22.

For some students, it was the creative aspects that engaged them. Some focused their efforts on drawing the art for their puppets or designing scenery to enhance their shows, like in Figure 4.19. Another Grade 2 student chose to make a 3D basketball court for his LeBron James paper puppet (see Figure 4.21). Others were drawn in by the script writing element; one pair of students, Rose and Bernadotte, decided as they worked on their script that their story needed "something evil." This led them to create a new puppet, a villainous troll swinging an animated spiked club (see Figure 4.20).

Figure 4.19: Grade 2 River stated during her interview: "I'm the background designer."

In the final sessions, Sonny encouraged each group to spend time putting the final touches on their creations. Some groups focused on adding to their script, some chose to spend time designing the set and props from cardboard pieces, and others wanted to perfect their character puppets with structural additions. The Grade 2 animatronics activities ended with a performance of their scripted shows in front of the class.

4.9 School Workshop Study Discussion

Our takeaways from the Grade 2 and Grade 6 study centered around the four major themes which emerged from our analysis process: creativity, challenge level, benefits of cross-grade mentoring, and suitability of the kit for elementary classrooms.

Combining Creativity and STEM with Puppet Design The focus on storytelling opens up endless creative possibilities. The parts in our kit were originally designed for the purpose of animating the mouths of talking characters. Perhaps unsurprisingly, the students immediately found creative ways to incorporate the simple linear motion to animate their characters, such as a pogo stick, a Whack-a-Mole game, and an ice cream cone. In the words of Grade 6 student Ryan, "There's pretty much no limits."

The open-endedness of the character design task provided student choice while allowing students to practice design thinking and problem solving. All the participants appreciated the creative freedom the kit afforded students. Both Anita and Sonny expressed in their post-interviews that students were highly engaged, and some threw themselves into the process of creating a character, writing a story, and designing and crafting the puppet. Students were able to bring their own interests into the stories, whether it be from popular culture or something more personal. Several Grade 2 students drew inspiration for their stories from video games like Minecraft (for example in Figure 4.18). Sonny remarked that "motivation comes from different places. It's really exciting for them to bring something like a character that they like to life."

Ryan's experience with STEM allowed him to understand the input and output of the Audio Board with no instruction, and independently create the bird in Figure 4.17. Charlie, who had knowledge of origami, showed his special friend how to make a 3D papercraft claw, which he actuated using the Linear Motor. But while these students were able to leverage knowledge from extra-curricular experiences, the majority of students, even in the Grade 6 class, struggled to find innovative applications for the kit, which could imply that creative mechanical design is not well scaffolded in the curriculum. Even using the components currently provided in our kit, it is possible to create more interesting motions, but few students possessed the engineering skills to experiment with them. Further research is needed to work towards equipping kids with these valuable skills. What is missing, perhaps, is a way to more creatively use the available technology – for example, using our limited motor mounts to create more interesting motions than 2D translation and rotation. The current curriculum doesn't seem to cultivate these skills, as evidenced by the similar complexities of the motions used in creations of the Grade 2s and Grade 6s.

Figure 4.21: 3D basketball court made from paper.

Students Becoming Teachers Cross-grade mentorship provides benefits for mentors and mentees, increasing engagement and providing students with opportunities for social-emotional growth. Anita spoke of the value of the creative partnerships in brainstorming and community-building, saying, "It was nice seeing them. The special friends were helping the bigger kids... I thought it was a really good relationship, bouncing ideas off one another." She described how the the Grade 2s provided direction on "how they would want the animatronics to work, like how fast, how slow. How it should move." Students also reflected positively on the experience. Sasha described seeing herself as a teacher: "It was kind of cool to hear myself explaining it to her ... because I hadn't really– I just kind of knew it in my mind, but it was cool to hear myself explain it." Thus, despite the lack of challenge felt by some of the Grade 6s, authentic motivation of creating a story for their special friend led to deeper and more sustained engagement.

Sonny reflected that having the experience with their special friends gave the Grade 2s a familiarity with the parts that allowed them to begin constructing characters independently right away. It also gave them examples of what a successful working puppet looked like, which guided their own design and building process. The collaboration with the Grade 6s provided direct inspiration

Figure 4.22: Grade 2 students perform their puppet show for the class.

for the Grade 2s' creations. For example, after seeing her special friend build a character swinging an axe with a Rotary Motor, Grade 2 student Fatima used a Rotary Motor to make her character swing a makeup brush in the same way. Another Grade 2 student expressed a very strong interest in using the Audio Board after seeing his special friend use it. Seeing their special friends work on animatronics also provided another source of motivation for Grade 2s, with the Grade 6s acting as role models. Sonny highlighted the intrinsic motivation for his students in working with older kids, saying, "I think that feels, you know, exciting for younger kids to feel like they're doing things that older students are doing."

Insights into Classroom Kit Evaluation The classroom environment is often chaotic and busy, and the classrooms in our study were no exception. One difficulty this introduced was the frustration students felt using the Mic Board, with many reporting that it didn't move the motors the way they expected it to. Students preferred boards that they felt gave them more control. Many students, especially in Grade 6, chose to use the Knob Boards or even the more complicated Audio Boards, instead of the Mic Board, indicating room for improvement on the technical implementation and interface of the mic board. The new PupCon board in Figure 4.5 was designed to give visual feedback while changing each setting using a knob. Perhaps in a future iteration, a push-to-talk button similar to walkie-talkies would alleviate the unwanted microphone response from other noises in the room.

In terms of character movement, some students encountered problems based on the materials they used to make the puppet. When students taped larger pieces of card stock paper to the zip tie, the zip tie would bend and twist, preventing the puppet from moving properly. In his post-interview Sonny said, "If it gets too big, it gets so floppy because like the motor's so small, there's not a huge backing to attach it to," demonstrating the need for a more rigid attachment than a zip tie. In one case, Sonny cleverly taped a popsicle stick to the zip tie, allowing the student's puppet to move the way she wanted.

Students also ran into technical hurdles with the motors. One type of servo developed a jamming issue, due to manufacturing errors, which occurred occasionally during our time with the Grade 2s, requiring a facilitator or student to gently pull on the zip tie to unjam it temporarily. When troubleshooting this and other technical issues, Sonny said that he became comfortable over time as he gained experience with the kit's components and materials. Anita, on the other hand, reported that she didn't have to do much troubleshooting at all with the Grade 6s, since they were self-sufficient when problem solving.

Challenge Level and Suitability for Elementary Students The kit provides a suitable challenge for younger elementary students but lacks the technical complexity to deeply engage older kids. One of our goals in providing easy-to-use PCBs specialised for animatronic puppetry is to lower the technical barrier to entry to begin telling stories, and the "plug-and-play" connectivity of components in our kit supported this. Anita said, "Once [the Grade 6s] had that first kind of lesson and their questions got answered, then they were off. So, it was a quick learning curve." Sonny told us "the benefit of having it so programmed is that it makes it really accessible."

The trade-off to having an easy-to-use kit made specifically for animatronics is the lack of open-endedness on the technical side. Students found that the fixed functionality of the motors and boards limits the complexity of possible creations. In our group discussions, the Grade 6 students expressed a desire for more complex types of motions and motor mounts, possibly akin to the gear systems used in [104], which would widen the design space. Anita also noted that the students "want to be more involved in the technology and the innovation of it" and suggested including details on the design of the PCBs themselves in future iterations of the kit. We were not able to include the Audio Board's feature of giving independent control of two motors because we could not install necessary software on the students' laptops. This could have provided a next step in the progression of difficulty through our kit.

In contrast with the older kids, Sonny's Grade 2 students were sufficiently challenged by the character construction. Mechanically planning the design proved difficult but in a good way. Sonny reflected, "That's awesome engineering problem solving for them. Trying to realize a character in those constraints is really good learning." Grade 2 students also struggled with motor skills required to physically build the character. River told us that cutting out the small pieces of the legs of her Minecraft creeper was the hardest part but was rewarding too. She said that her favourite part was "seeing what it would look like when [she] was done."

4.10 Robotics Camp Workshop 2

After the school workshop study, we again collaborated with the robotics institute from Section 4.6 who ran another 1-day workshop at the end of a 1-week robotics camp. Thirty-four students aged 7-16 attended.

Our goal this time was to 1) improve on the aforementioned aspects of classroom management and age range issues from the last workshop and 2) try to incorporate programming into the tasks using the Audio Board, Arduinos, and the Servo Shield. As discussed in Section 4.4 about the kit, our Audio Board can be used to trigger events with an Arduino microcontroller.

From the classroom management issues we encountered during the first camp pilot study, we better scaffolded the day's activities thanks to the robotics institute staff, who created an online instructable that students could work from at their own pace (see Figure 4.23).

Students were also arranged in groups at tables with peers around the same age, which allowed for tables with younger students to spend more time on the tasks they struggled with and for teachers to answer questions for the whole table all at once. This eliminated the waiting period where students would raise their hands for help but all staff members and research team members were busy. Older students were able to move onto the more advanced tasks for the day by going through the module on their laptops.

Like the first robotics camp workshop, the morning activity was to create a puppet either from a given template or to make their own character. Again, most students chose to use a template but some chose to decorate the template to give it their own personal flair. For example, two girls who used template animal puppets showed their creativity. One student added hair to the penguin and the other added earrings to the zebra, modeled after members of a K-pop band they enjoyed.

All students successfully completed the morning activity. After lunch, students moved on to making 2 character shows with the Audio Board, and then further onto a programming task. The programming task involved connecting jumper wires from the Audio Board to an Arduino and also plugging in the servo shield from our kit to the Arduino. Using skeleton code and guiding themselves through the online module, the students connected three linear motors to the servo shield and programmed simple movements in the Arduino IDE to animate buzzing bees. About half of the students reached this portion of the assignment, and around 10 students finished up the task right as the day ended and all the equipment had to be taken apart and given back.

Overall, this robotics camp workshop was much more successful than the first one in terms of student progress and classroom management.

4.11 Broader Discussion Including Robotics Camps

Age Group Considerations Dealing with different age levels can be difficult as we've seen in our school user study, and this is exacerbated when the age range is even bigger. The younger students can struggle with cutting out their characters and putting together the motors, while the older students finish this task quickly, and if the task is improperly planned, have nothing to work on next. In the first robotics camp workshop, this led to a chaotic and frustrating classroom environment. Thus, we tried to give the older students programming tasks to work through independently using prepared online slides on their own laptops in the second workshop, like those in Figure 4.23. This worked very well because this time around, students were split into age groups at each table. That way, the tables with the younger groups got to work on the skills at their level and help each other out, and the tables with older groups were able to move onto the programming portion of the task at their own pace. The first robotics camp workshop let students sit with their friends and siblings which, while fun for the students, led to a mismatch of skill levels. The classroom management, although lively with all the students working on their animatronics, was much calmer and more organised than the first workshop.

Student Backgrounds It's worth noting that the students who took part in this workshop were self-selected for doing technical activities, in that they were already participating in a robotics

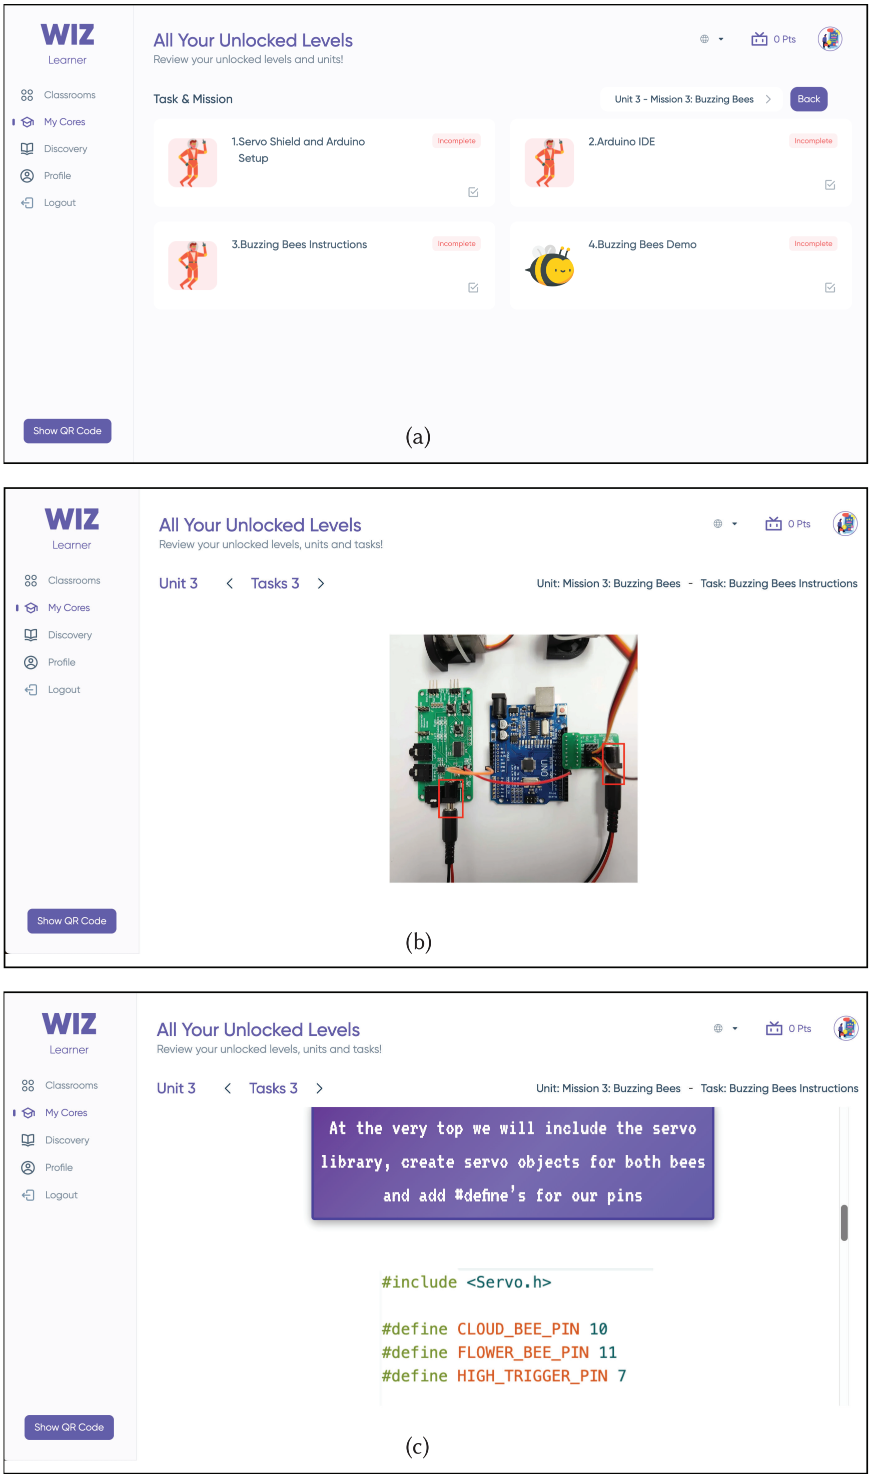

Figure 4.23: The module made by the robotics institute staff. The Animatronics module on the online platform (a) has 4 sections. The first section (b) shows students how to plug in the Servo Shield and Audio Board into the Arduino. A section later in the activity (c) explains the Arduino code step by step to help students if they get stuck.

summer camp. Many of them had previous programming and robotics experience before starting our animatronic workshop. Naturally, these kids are either self-motivated to learn technical skills or their parents had signed them up to participate. Parents are often eager and perhaps a bit pressured to have their children learn to code, as it has become a desirable skill in society.

The K-6 students in our school study did not have a choice to participate since the assignment was done during their regular class time. In a more general student population, students will have a wide range of previous experience with technology and programming, and as noted earlier, many of them preferred the craft and storytelling aspect of the animatronic creation process. However, in the robotics camp workshops, we observed that very few students would have chosen to do the art portion of the activity themselves as evidenced by the number of students who chose to use given template puppets in our kit rather than create their own character. The second robotics camp workshop had a higher level of technical difficulty than what we did in the school, which was expected due to its specialised nature. To be fair, there is not enough time to teach the students the basics of storytelling, have them come up with original characters and dialogues, craft the character, construct the puppet, and get to the programming stage in just 1 day. In the future, extending the robotics camp workshops to two or even three days would help us blend the technical and storytelling aspects together better.

Technical Teachers vs. Non-technical Teachers The robotics camp workshops and the pilot study with the JK students were all run by technically proficient teachers with plenty of personal and teaching experience in the field. They were able to independently plan their workshops with little input from our research team other than briefly showing them how the parts in the kit worked. The second robotics camp workshop featured a nicely planned and aesthetically pleasing online module about animatronics that their staff made the week of the workshop. The module worked extremely well in this context, but it is fair to say that non-technical teachers have neither the time nor training to create their own animatronics curriculum. Developing more teacher training tools is necessary if we want using animatronics in the classroom to be a feasible option for teachers without a strong technical background.

Even without a tech background, the Grade 2 teacher was able to troubleshoot for himself pretty well. We saw his clever problem solving skills throughout the study. The Grade 6 teacher's students were able to work out problems amongst themselves for the most part. However, in both classrooms the research team had to jump in to diagnose in-the-moment issues with the hardware in order to keep the flow of progress.

4.12 Storytelling + Programming with Audio Boards

The versatility of animatronic activities and utility of our kit is difficult to show accurately, with most students we worked with being Grade 6 or below and with such a small sample size. There are many other types of possible projects to build with the kit, which we expect would challenge middle and high school aged students and give them the right amount of difficulty and expressivity.

As mentioned in our study, the Grade 6 students were hungry for more technical challenges after successfully using the two easier boards in their animatronic stories. The Grade 6 teacher in her post-interview reiterated to us that her students craved to move onto the next more technically difficult steps of animatronic activities using the Audio boards. Understandably, high school students would most likely also become bored with the simple live puppeteering tasks we did with the younger students.

One of our goals with bringing animatronics into the classroom is to combine the technical skills with storytelling in a compelling way to show *all* students that they are capable of *every* part that goes into making an animatronic show. The following example animatronic shows demonstrate the extra capabilities of the Audio Boards which would typically be done by middle and high school aged students, although some of the Grade 6 students expressed interest in doing more advanced animatronics shows.

Using the Audio Board's feature to control two motors independently, students can record conversations between two characters on the right and left tracks of an audio file to make the animatronic characters talk to each other, as seen in Fig 4.24. Two meerkats speak to each other about their nocturnal nature, educating viewers about the type of animal they are.

Figure 4.24: Two meerkats discuss their sleeping habits.

As mentioned, the Audio Board can connect to an Arduino to add programmable cues into the show. Seen in Figure 4.25 the Elk example utilizes the LED shield, an attachment in our kit that plugs into an Arduino microcontroller, to trigger LED lights to aid the story. The elk educates viewers about different holidays, lighting up the LEDs as they talk about each one.

Figure 4.25: The elk teaches the audience about Christmas, Hanukkah, Kwanza, Diwali, and New Years. The LED shield attaches to the Arduino and allows easier programmatic control during the show.

In the example in Fig 4.26, dialogue on the right audio track makes the character talk, and the left audio track can be used to control other motors by programming events to happen at each beep in the track. The Billy Bear example contains a talking bear telling the audience about the role of bees in the environment, as the bees constantly buzz throughout the show. When Billy describes that bees are in danger, they are cued to “disappear” behind the clouds, using Linear Motor units attached to the zip ties. Billy’s eyes also move back and forth during his talking, giving a more realistic performance. The servos used in this show take advantage of the servo shield in our kit, pictured in Fig 4.7, to control five servos at once.

Figure 4.26: Billy Bear explains the importance of buzzing bees around him.

An even more advanced assignment can add interactivity into the animatronic show. The following example in Fig 4.27 is intended for upper year high school students to do over a longer period of time. The interactive Shakespeare example not only includes written dialogue, intermittent eye blinking, and glancing but also is connected to a backend written in Python which communicates with the ChatGPT API. Users can speak into the microphone and ask Shakespeare a question. The audio is converted from speech to text, which goes into a ChatGPT prompt asking the large language model to respond to the question in the style of Shakespeare. The response from ChatGPT is converted from text to speech and said by the animatronic character using our Audio Board. This can generalise to other historical figures, famous people, or original characters; it simply requires changing the ChatGPT prompt. Because ChatGPT can take a few seconds to respond, we have included pre-written lines of audio in the style of Shakespeare which are played while the program waits for the returned ChatGPT response. Additionally, if an assignment like this is used in the classroom, students can analyse whether the language model produced an accurate response, providing another layer of learning the topic.

4.13 Future Work: Impact of Animatronics in K-12 Classrooms

We presented a Paper Animatronics Kit aimed at K-12 students along with four separate pilot studies. Our two Robotics Camp 1-day Workshops shed light on the experience and limits of working with large age ranges of students in a very short amount of time and allowed us to quickly test new kit components. We also conducted a School Workshop user study in a K-6 lab school to validate its suitability in a classroom context, including a pilot study with Junior Kindergarten students and a longer term study with two older grade levels. Constructing puppets with JK students allowed for initial observation of the kit and animatronic building process with very young students and prepared us for our largest user study in this thesis. Working with the Grade 2 and Grade 6 teachers and students, we aimed to evaluate the benefits and challenges of using animatronics with our kit in

Figure 4.27: ChatGPT Shakespeare can answer questions in old English.

their classrooms. Our results indicate that the kit was effective at engaging students in the creative process, and provided opportunities for cross-curricular integration of STEM and literacy.

We believe that the interdisciplinary nature of animatronics provides an effective way of bridging creative and technical activities by providing multiple entry-points for varying student interests. Incorporating cross-grade mentoring gave students motivation and inspiration to tell stories, fostered collaboration and idea-sharing, and encouraged self-reflection of prior learning.

We are interested in improving the interfaces and functionality of the boards in our kit. For example, being able to program motions with the Knob Board would give students more control. Additionally, more investigation is needed into the state of creative and design thinking skills in STEM education to uncover how best to develop and scaffold these skills across the curriculum.

In our initial conversations with Grade 2 teacher Sonny, we discussed using animatronics in their science unit learning the life cycle of the Atlantic salmon. Every year, the Grade 2 classroom hosts an aquarium and habitat for salmon, watching them go from salmon eggs all the way to releasing them in the wild. The activities would be salmon-themed, and the class would split into groups, each studying a different phase of life. They would create salmon egg and fish puppets of all backgrounds telling the stories of their characters while learning important science concepts. Because of time constraints and for ease of fitting animatronics into the classroom, Sonny chose to more seamlessly fit the animatronics into the Grade 2 literacy curriculum. We want to explore more cases of using this medium to learn subjects *other than* literacy – rather than just augmenting creative writing with animatronics – and are interested in utilizing storytelling to teach students science, history, and more.

A more long term goal with animatronics in the classroom is to empower the students using the kit not only to express themselves through this art form but also to challenge the way they see themselves and their skills. An 8-week study is not enough time to examine the long term impact of teaching kids that they can be technically inclined *and* be creative, not one or the other. We observed students who prefer crafting their characters successfully understand the mechanical concepts behind making it move with the boards. We also saw students who told us they only liked the "tech" part make meaningful characters and engage with their classmates. This is a promising step in the right direction, but further study is required.

4.14 Appendix: Interview Questions

Teacher Semi-Structured Pre-Interviews The following interview questions were a way for us to explore with teachers what they would want to get out of running an animatronics workshop in their classroom. We felt this step was important, as the whole point of the collaboration was to create a mutually beneficial user study where us researchers could figure out how best to support educators in integrating our animatronics kit into their curriculum. Teachers we interviewed were not obligated to then run a proposed workshop; these questions were purely hypothetical. However, two of the teachers we interviewed agreed to participate in the next phase of the study, and the results of the semi-structured interviews helped influence and inform the workshop we ended up doing in the Grade 2 and Grade 6 classrooms.

Research Questions for Teacher Interviews for Animatronics in Classrooms

- What type of workshops (using our kit) would they be interested in running?
- What do they hope to teach their students through an exercise like this?
- What are they interested in learning for themselves? e.g. some of the capabilities of technology
- What level of expertise is needed from a teacher to run an animatronics unit in their classroom?

1. Background/History of Teacher and Area of Expertise
    - Education history: what degrees, what level of education, what certificates if any do you have
    - Teaching history: where have you taught, for how long, what role, what age groups
    - Expertise: specialties, concentrations, other skills and experience
2. Subjects
    - What subject(s) do you teach?
    - What made you want to teach that subject? (if they are a subject-specific teacher)
    - What is your favorite subject to teach? (if they are a primary school teacher who teach all subjects)
    - How much freedom do you have to implement new workshops into your classroom?
3. Student Demographics
    - What age groups do you teach?
    - What gender groups do you teach?
    - What income level groups do you teach?
    - Do you see clear differences between how these groups learn? react to assessments? engage in the classroom?
    - Do you teach students with disabilities?

- How do you approach teaching or engaging a student who dislikes a certain subject with various class activities or lessons? (e.g. a student doesn't like math so it's hard to get them interested or feeling good about doing math problems in class)

4. Comfort with Technology

   - Have you used a robotics or electronics hardware kit in the classroom before?
   - What was it called?
   - What did you use it for?
   - Did you like it?
   - Did your students like it?
   - Would you feel comfortable using this kit in your classroom without help? (After appropriate training and practice?)
   - Why or why not?
   - What parts do you find confusing?
   - What parts do you find intimidating?
   - What kind of accessibility features would be needed in order to make the animatronics workshop accessible?

5. Evaluation

   - How to assess if students learned better or engaged more in the subject using the kit?
   - How do you normally get a feel for how well a student learns or how excited they are about a topic?
   - Do you expect the students will enjoy completing the assignment you designed?
   - How to evaluate how well a story was conveyed?

Grade 6 Post Exploratory Learning Group Discussion After the first session of the animatronics workshop with the Grade 6 students, we led a 10 minute group discussion with the whole class. The Grade 6 students had 1 hour to figure out on their own how to use two of the boards in our kit. The guiding group discussion questions were as follows:

- Please describe what you made. Who is your character?
- Why did you choose that character?
- How did you come up with the script?
- Did you have fun making the puppets?
- What did you find enjoyable about the process?
- What challenges did you run into during the process?
- Ask the student to walk through their problem solving process during the challenge.
- What other things did you want to make your puppet do that you couldn't?

Grade 6 Post Mentoring Group Discussion After the first Special Friends session where the Grade 6 students made puppets with their Grade 2 assigned student pairs, we led another 10 minute group discussion with the Grade 6 students once their special friends returned to their own classroom. We wanted to ask the Grade 6 students if their own understanding of the boards changed after teaching it to their younger friends. The guiding group discussion questions were as follows. During discussion we added follow up questions depending on how the conversation went.

- Which boards did you use for your special friends and why?
- When you were working with them, did you try to encourage them to figure out the moving parts? (Or did you do it for them?)
- Did you explain the parts to them? Or did you demonstrate and show them?
- Last time I was here, you all did really well and figured out how they work basically on your own. So when you were teaching your special friend, do you think it gave you a better understanding of the boards, or you felt like you already knew how to do it?
- Which new things did you learn about the functionality, if any?
- In teaching the board to your special friend, do you feel like that enhanced your learning? Or do you feel like it was the same?

Grade 6 Individual and Small Group Interview Questions After the Grade 6 students spent more time on the shows that they were writing and presenting to their special friends, we interviewed individual students or groups of 2-3 students at a time. We asked them about

- What's your story?
- Ask follow up questions about their story.
- How did you come up with story and the design of your character?
- How does it move?
- So you knew it was going to move when you were making it, right? So did that affect the drawing part? Did you think about it while you were drawing?
- What did you find difficult about it?
- If you're working with a partner or group, how did you split the work? Are you working on the art or tech part?
- Did you have a favourite part about the activity? What did you like more (art part or tech part)?
- Name another thing you might want to make with animatronics.
- What other school assignment would you want to do (or re-do) using animatronics?
- Would you want more assignments like this?
- Do you think of animatronics as some way to communicate or is it like more creative storytelling?

Grade 2 Individual and Small Group Interview Questions We also interviewed the Grade 2 students individually or in small groups after they worked on their animatronic shows. The Grade 2 students had each been writing their own creative stories, and chose a character to make a puppet by themselves (after having done it with their special friend). Then they formed groups and wrote scripts involving their characters meeting each other. We interviewed some groups of the Grade 2 students. The questions were similar to the Grade 6 interview questions. We also asked about their special friend dynamics since we had only talked to the Grade 6 students about the mentoring aspect at that point. We asked a subset of the following questions depending on the patience of the Grade 2 students, the flow of conversation, and trying to gauge how well the student understood the question worded a certain way.

- Tell me about your character.
- What's your story?
- How did you decide to make that specific character from your story instead of any other?
- Did you find any part of this process hard? Which part?
- Are you happy with what you made?
- Did you find it difficult?
- What was your favourite part?
- Did you learn a lot from your special friend (or did you figure it out by yourself)?
- Did you like writing the story or building the animatronic?
- How did you make it?
- Do you think building it with your special friend helped you build it yourself?

Teacher Semi-Structured Post-Interviews After the school study was completed, we went back to talk to the Grade 2 and Grade 6 teachers about how they thought the animatronic activities went based on their observations. To help answer our research questions, we wanted the teachers' perspective on student engagement, mentoring, curriculum integration, and the kit itself.

1. Student Engagement
    - Comment on the combination of art and STEM.
    - Did you notice any students who seemed more into the activity than usual? (focus and motivation)
    - (for me) Restate goal of activity: to create entry points into art/storytelling and STEM for students who prefer one or the other. (for the teacher) Do you think we motivated STEM learning through storytelling?
2. Cross-Curricular Integration
    - Were you satisfied with how their stories/puppets turned out? Did anything surprise you?

- How do animatronics fit into your curriculum goals? If you could fit animatronics into another unit, how would you do it?

3. Mentoring

- Do you think having the students teach their special friend strengthened their own understanding?
- Did it affect what they made?

4. Teacher Comfort

- How did you find facilitating something like this? i.e. complexity/equipment/busyness/classroom management
- How did you find implementing a STEM activity?
- How comfortable would you be doing this without us and/or the tech teacher – fewer adults in the room?

5. Kit

- What other parts do you think would be useful to add to the kit?
- Which things could be easier? What did they struggle with?

CHAPTER 5

Conclusion

The main question of this thesis was: How can physical storytelling be enhanced by improving each of the three mediums: floating sculptures, zoetropes, and animatronics? The innovations presented here are not only technical advancements; they redefine how audiences and creators interact with these forms. Conceptually, the enhancements of these mediums lead to a more immersive and accessible experience and allow for deeper storytelling than before. The mediums discussed throughout this thesis certainly do not cover the entire field of physical storytelling, but instead form a representative set that spans a wide range of experiences and challenges within physical storytelling. Each medium — floating sculptures, zoetropes, and animatronics — represents a different approach to immersing audiences physically in storytelling. By addressing the structural challenges of floating sculptures, the narrative limitations of zoetropes, and the accessibility barriers in animatronics, this thesis has expanded what is possible in terms of both the depth and breadth of physical storytelling.

Sculptures represent the long tradition of human expression through still objects like cave drawings, paintings, and ceramics to name a few. Still objects capture a single moment in time; a physical space with arranged sculptures can immerse viewers in the space with the scene. A challenge central to creating a space full of sculptures has been how and where to mount the objects such that they tell the story the artist is creating.

In the case of ***floating sculptures***, I explored how to adapt mechanical engineering optimization techniques to balance force and torque constraints while maintaining an invisible support structure. By developing an algorithm that conceals these supports, I contributed a novel approach to creating immersive, walk-through experiences where the audience can engage directly with the suspended sculptures. This technique not only opens up new possibilities for gallery exhibits but also empowers creators to design intricate, floating environments that feel magical and structurally sound. The validation of this method through real-life wire and rod structures demonstrated its feasibility and practical potential in physical storytelling.

Floating sculptures now invite audiences to step into a scene that feels truly unsupported and untethered, creating a new layer of immersion. Immersion is an important aspect of any storytelling, and physical stories are no exception. Suspension of disbelief is the key to experiencing and enjoying a fictional story. Any visible rods can distract the viewer and take away from the narrative being told with the objects and 3D space of the exhibit. Additional improvements to the method, such as allowing rods to bend under load instead of requiring straight truss elements, would enable more complex and thin objects to be supported invisibly. The tradeoff here is that this inevitably complicates the optimization problem by introducing non-linearity which would cause the solver to take longer to find an optimal solution.

Another challenge to consider is that the shadows of the rods in my resulting structures may be visible, even if the rods are not. My method as it stands does not consider light sources in the scene as part of the visibility considerations. A light source in the scene could be treated as another viewer (or viewpoint distribution) looking into the scene, and this could be easily added to the computation of the visibility cost for each rod using standard ray tracing techniques.

In terms of artist experience, my method could benefit from a more visual way to instruct the maker in assembling the scenes. For the flying seagull example in Chapter 2 Figure 2.14, I needed to know the precise locations of the wires' attachment points to the supporting surface and scene objects and a way to attach them. This required adding small torus shaped meshes to act as hooks which I could tie fishing wire to, and it required exporting the meshes of the scene along with the hooks and importing the geometry to 3D printing software. Similarly, assembling scenes with rods was equally painful as I needed to carve indents or holes for the rods into the meshes of the scene at attachment points. Constructing arrangements of floating sculptures through a usable graphical interface rather than running scripts would be preferable and allow for more creators to make scenes using my method.

Applying this method to create a larger scale storytelling walkthrough would be exciting and further prove the viability of the method. One idea for a large scale exhibit is a spooky children's playroom. This would involve purchasing objects such as toys to be levitating around the room and then creating a digital version of the scene to be run through the algorithm. The viewpoints would follow a path through the 3D space rather than a distribution on a plane. The immersive experience of traversing the playroom would reveal that the dolls and horses are haunted through audio and potentially even LEDs in the dolls' eyes.

Large scale fabrication is an interesting subfield on its own, and many of the techniques could be useful when it comes to creating a large exhibit. For example, building information modeling (BIM), is a methodology used in architectural projects to plan and visualise buildings before they are constructed and to aid the construction process through software [6]. Baudisch et al. similarly have made large scale truss fabrication hardware and sofware systems which aim to let users create truss structures from steel and springs [76] or automatically from inflatable material [118].

Zoetropes take still sculptures and add life to them with movement through time. Zoetropes are the predecessor to modern film, and the medium quickly evolved into telling longer form stories projected in movie theatres. The advantage to zoetropes over film is the ability to watch truly 3D movies rather than 3D scenes projected onto images. As discussed in Chapter 3, a drawback of other 3D displays is that projections and screens are backlit, removing any way to interact with the scenes using light. Hirsch et al. propose a bi-directional light field display to mitigate these issues and maintain a glasses and headset free experience, a feature my zoetrope also has [56].

To work within the short and periodic limitations of zoetropes and take advantage of interaction with light, I expanded on the traditional medium by incorporating audio triggered by user-directed light, enhancing both the interaction and narrative depth of the stories they tell. My work revealed hidden elements within a scene, allowing viewers to experience plot twists and uncover surprises through active participation. By giving the audience the ability to trigger audio cues, I pushed the boundaries of how a zoetrope can engage viewers and challenged their assumptions, making the story more immersive and multifaceted. This innovation allows for more complex storytelling within the limited physical constraints of zoetropes.

There is ample room for injecting more interactivity into zoetropes. Along with visual and audio cues to tell the story, other senses could be engaged; for example a story taking place outside during a chaotic storm could feature blowing wind and the smell of rain to further immerse the viewer. Additionally, more information can be packed into my 3D zoetrope and revealed in different ways. Taking advantage of the strobe light, certain objects in the scene can be hidden until viewed with UV light.

An artist could interleave stories between even and odd frame sets within the same zoetrope, switching between them when the user focuses on different parts of the scene or after a set amount of viewed revolutions of the wheel. One scene could be displayed on every even-numbered frame in the zoetrope, flashing the light only for those frames. If the viewer shines the light on a specific part of the scene, the zoetrope could switch to flashing the light on only odd frames to reveal a slightly different scene, subverting initial expectations of the story.

The Eigen Zoetrope introduces the idea of piecing together a single frame from multiple frames, adding together parts of each frame and fusing them using very fast motion and light [75]. An example that could be made using this technique is a story involving ghosts. By spinning the zoetrope so quickly that flashing the light on multiple frames seemingly in the same place, an artist could create the illusion of transparency. The two frames could contain the background of the scene and only one of those frames could contain a ghostly character. When viewed at almost the exact same time, both frames combine to make a solid background with a transparent ghost. This could also be a clever way of hiding support structures for floating objects!

It would be really fascinating to be able to generate 3D scenes that look the same upside down so that a viewer on each side of the wheel could experience the story instead of just one at a time. More complex is the idea of having two separate stories for the viewers on the two sides, where each frame must be semantically meaningful *and* temporally coherent. Using diffusion models, Geng et al. were able to produce multi-view optical illusions in the form of images displaying one scene when viewed at first and a different scene when the image is flipped [44]. Temporal coherence and an extension of their method into 3D could enable zoetrope anagrams.

To physically fit more frames of animation into the story, a spiral shaped zoetrope or a linear zoetrope could allow for longer or more detailed stories to be told, although this is a grand engineering feat.

Another avenue for future research could address the inaccessibility of zoetropes to people who cannot look at flashing lights. Currently that population is completely unable to view my zoetrope animations in person. One way to approach this problem is by designing a ratcheting wheel spinning mechanism, similar to the way film projectors work. Reaching more people with stories in this interesting format is crucial to the idea that storytelling is about community, empathy, and self-expression. Anyone should be able to tell a physical story with these mediums in an ideal world.

Animatronics also move sculptures, but instead of showing still frames consecutively, they show a single frame of moving sculptures. The roots of animatronics came from mechanically moving automata all the way to large scale physical talking and moving lifelike robot characters we see today in theme parks.

With the medium of animatronics, I focused on making this traditionally complex medium more accessible, particularly within a K-12 educational context. By developing an affordable, versatile kit that blends papercraft and simple electronics, I offered students the tools to tell stories through physical puppetry. My approach showed that animatronics could democratize storytelling for students, regardless of their technical background, encouraging them to explore mechanical creation, character design, writing, and performing. Through this work, I challenged the perception that students must be either artists or engineers, helping them see their potential as creators across disciplines.

My work in paper animatronics can be expanded further as well. New and different types of servo motor mounts would let students make more complex characters that move in ways other than just rotary and linearly with joints that have more degrees of freedom. The paper characters could also move between scenes, possibly using wheels or following along a magnetic track. It could push students to come up with more complicated stories, for example, with moving set pieces. Even more lifelike, it would add to the story if the animatronic puppet's mouth shapes more closely matched the input audio and mimicked human facial expressions during speech.

Students could create interactive animatronic shows along the lines of the Shakespeare puppet that takes questions from the audience, mentioned in Chapter 4. This would require more tools in the kit for both students – who would need to write code, and teachers – who would have to integrate this into their curriculum.

Forms of interaction other than voice are also an exciting step forward. Using buttons, sensors, and cameras for other user input to affect the animatronic display can have a big impact on the immersion, participation, and interest in a show. Of course, making more complex shows requires intuitive tools for young students to be able to create them. Adding complexity into the stories would most likely be intimidating for teachers in terms of classroom time and technical skills. Especially for teachers without technical training, this issue must also be addressed. Potentially a booklet or guide containing past project examples and a framework of how to design an animatronics lesson plan could be developed and provided to the teachers.

Reaching students of less privileged socio-economic status is a necessary step forward in the field of animatronics (but more broadly tangible storytelling mediums) in the context of education. Indigenous Canadian students of all ages, but particularly high school aged students, rarely get to experience that level of education at all, much less using traditionally expensive robot components. Their stories are extremely important to hear and be told by them through their own lens of the world because they have historically been silenced and ignored.

Together, the contributions of my thesis provide a foundation for further exploration in physical storytelling, encouraging artists, engineers, and educators alike to build on this work and push its boundaries even further. By bridging disciplines and expanding the tools available to creators, this research contributes not only to the development of these specific media but also to the larger landscape of physical storytelling, making it richer, more inclusive, and more engaging for a diverse range of participants and audiences.

The future of physical storytelling is bright. Floating sculptures, zoetropes, and animatronics can be experienced in new ways by audiences, from classrooms to galleries and even theme parks. Physical storytelling will continue to grow, as these mediums become more accessible and integrated into both educational and entertainment spaces. The possibility for more people to not only experience but also tell their own stories is greater than ever before, making storytelling a truly universal form of expression.

Bibliography

[1] Adobe. *Mixamo*. 2024. URL: https://www.mixamo.com/.

[2] Jennifer Ginger Alford, Lucas Jacob, and Paul Dietz. "Animatronics Workshop: A Theater + Engineering Collaboration at a High School". In: *IEEE Computer Graphics and Applications* 33.6 (2013), pp. 9–13.

[3] E. D. Andersen and K. D. Andersen. "The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm". In: *High Performance Optimization*. 2000.

[4] Michelle Annett et al. "MoveableMaker: facilitating the design, generation, and assembly of moveable papercraft". In: *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. 2015, pp. 565–574.

[5] Autodesk. *Fusion 360*. Autodesk. 2024. URL: https://www.autodesk.com/ca-en/products/fusion-360.

[6] Salman Azhar, Malik Khalfan, and Tayyab Maqsood. "Building information modeling (BIM): now and beyond". In: *Australasian Journal of Construction Economics and Building, The* 12.4 (2012), pp. 15–28.

[7] Moritz Bächer et al. "Spin-it: optimizing moment of inertia for spinnable objects". In: *ACM Trans. Graph.* (2014).

[8] Reynold Bailey et al. "Subtle gaze direction". In: *ACM Transactions on Graphics (TOG)* 28.4 (2009), pp. 1–14.

[9] Felix Barrett and Maxine Doyle. *Sleep No More*. 2011.

[10] Gregory Barsamian. *Feral Fount*. 1996.

[11] Sally Barton-Arwood, Kristine Jolivette, and N. Gayle Massey. "Mentoring with Elementary-Age Students". In: *Intervention in School and Clinic* 36.1 (2000), pp. 36–39.

[12] Ayah Bdeir. "Electronics as material: littleBits". In: *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*. Association for Computing Machinery, 2009, pp. 397–400.

[13] G. Bradski. "The OpenCV Library". In: *Dr. Dobb's Journal of Software Tools* (2000).

[14] Jean-Louis Brenninkmeijer. *Little Canada*. 2021.

[15] The Editors of Encyclopaedia. Britannica. *Kinetoscope*. 2020. URL: https://www.britannica.com/technology/Kinetoscope.

[16] Ian Brodie. "Stand-up Comedy as a Genre of Intimacy". In: *Ethnologies* 30.2 (2008), pp. 153–180.

[17] Kurt A Bruder and Ozum Ucok. "Interactive art interpretation: How viewers make sense of paintings in conversation". In: *Symbolic Interaction* 23.4 (2000), pp. 337–358.

[18] Roberto Brunelli. *Template matching techniques in computer vision: theory and practice*. John Wiley & Sons, 2009.

[19] Alexander Calder. *Romulus and Remus*. Medium: Wood, steel wire, and springs. 1928.

[20] Emmanuel J Candes, Michael B Wakin, and Stephen P Boyd. "Enhancing sparsity by reweighted l1 minimization". In: *Journal of Fourier analysis and applications* (2008).

[21] Connie Champlin. *Storytelling with puppets*. American Library Association, 1998.

[22] Subramanian Chidambaram et al. "Shape Structuralizer: Design, Fabrication, and User-driven Iterative Refinement of 3D Mesh Models". In: *Proc. CHI*. 2019.

[23] Brandy Clark. *'Avatar' Re-Release Sees 93% of Domestic Audience Choose 3D for Their Return to Pandora*. 2022. URL: https://collider.com/avatar-re-release-3d-viewership-93-percent/.

[24] Alvin P Cohen. "Documentation Relating to the Origins of the Chinese Shadow-Puppet Theater". In: *Asia Major* (2000), pp. 83–108.

[25] Fiona Collins. "The use of traditional storytelling in education to the learning of literacy skills". In: *Early Child Development and Care* 152.1 (1999), pp. 77–108.

[26] Blender Online Community. *The Free and Open Source 3D Creation Suite*. Blender Foundation. Stichting Blender Foundation, Amsterdam, 2024. URL: http://www.blender.org.

[27] Robert Connelly and Walter Whiteley. "Second-Order Rigidity and Prestress Stability for Tensegrity Frameworks". In: *SIAM J. Discrete Math.* (1996).

[28] Stelian Coros, Jonas Zehnder, and Bernhard Thomaszewski. "Designing structurally-sound ornamental curve networks". In: *ACM Trans. Graph.* (2016).

[29] Donald Crafton. *The talkies: American cinema's transition to sound, 1926-1931*. Vol. 4. Univ of California Press, 1999.

[30] Arthur J Critchlow. *Introduction to robotics*. MacMillan Press Ltd., London, England, 1985.

[31] Jennifer L. Cross et al. "Student outcomes from the evaluation of a transdisciplinary middle school robotics program". In: *2017 IEEE Frontiers in Education Conference (FIE)*. 2017, pp. 1–9.

[32] Mario Deuss et al. "Assembling self-supporting structures". In: *ACM Trans. Graph.* (2014).

[33] Michael Dieter and Geert Lovnik. "Theses on Making in A Digital Age". In: *Critical Making - Manifestos*. Ed. by Garnet Hertz. Telharmonium Press, 2012.

[34] Harish Doraiswamy et al. "Topology-based catalogue exploration framework for identifying view-enhanced tower designs". In: *ACM Trans. Graph.* (2015).

[35] W Dorn. "Automatic design of optimal structures". In: *J. de Mecanique* (1964).

[36] Kenneth Elpus. "Access to arts education in America: The availability of visual art, music, dance, and theater courses in US high schools". In: *Arts Education Policy Review* 123.2 (2022), pp. 50–69.

[37] Shuangkang Fang et al. "Chat-edit-3d: Interactive 3d scene editing via text prompts". In: *European Conference on Computer Vision*. Springer. 2025, pp. 199–216.

[38] Mingbin Feng et al. "Complementarity formulations of l0-norm optimization problems". In: *Industrial Engineering and Management Sciences. Technical Report. Northwestern University, Evanston, IL, USA* (2013).

[39] Robert S Fisher et al. "Visually sensitive seizures: An updated review by the Epilepsy Foundation". In: *Epilepsia* 63.4 (2022), pp. 739–768.

[40] Meow Wolf Foundation. *Omega Mart*. 2021.

[41] Laccone Francesco et al. "Automatic Design of Cable-Tensioned Glass Shells". In: *Comput. Graph. Forum* (2020).

[42] Robert M. Freund. "Truss Design and Convex Optimization". PhD thesis. M.I.T., 2004.

[43] André Gaudreault, Nicolas Dulac, and Santiago Hidalgo. *A companion to early cinema*. John Wiley & Sons, 2012.

[44] Daniel Geng, Inbum Park, and Andrew Owens. "Visual anagrams: Generating multi-view optical illusions with diffusion models". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024, pp. 24154–24163.

[45] Karine Giboulo. *Housewarming*. 2023.

[46] Keith Gluck. "The early days of audio-animatronics". In: *The Walt Disney Family Museum Blog* (2013).

[47] Thomas Görne. "The emotional impact of sound: A short theory of film sound design". In: *EPiC Series in Technology* 1 (2019), pp. 17–30.

[48] Akinori Goto. *CROSSING #03*. 2019.

[49] Tovi Grossman, Daniel Wigdor, and Ravin Balakrishnan. "Multi-finger gestural interaction with 3d volumetric displays". In: *Proceedings of the 17th annual ACM symposium on User interface software and technology*. 2004, pp. 61–70.

[50] Emily Hamner et al. "Robot Diaries: Broadening Participation in the Computer Science Pipeline through Social Technical Exploration." In: *AAAI spring symposium: using AI to motivate greater participation in computer science*. 2008, pp. 38–43.

[51] Emily Hamner et al. "Training teachers to integrate engineering into non-technical middle school curriculum". In: *2016 IEEE Frontiers in Education Conference (FIE)*. 2016, pp. 1–9.

[52] Nazanin Sadat Hashemi et al. *Template Matching Advances and Applications in Image Analysis*. 2016. arXiv: 1610.07231 [cs.CV]. URL: https://arxiv.org/abs/1610.07231.

[53] Tracy Ann Hayes, Theresa Edlmann, and Laurinda Brown. *Storytelling: global reflections on narrative*. Vol. 122. Brill, 2019.

[54] Fritz Heider and Marianne Simmel. "An Experimental Study of Apparent Behavior". In: *The American Journal of Psychology* 57.2 (1944), pp. 243–259.

[55] Garnet Hertz. "Interview with Matt Ratto". In: *Critical Making - Conversations*. Ed. by Garnet Hertz. Telharmonium Press, 2012.

[56] Matthew Hirsch et al. "8D: interacting with a relightable glasses-free 3D display". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '13. Paris, France: Association for Computing Machinery, 2013, pp. 2209–2212.

[57] Atlas Obscura HJHausman. *Masstransiscope: A modern zoetrope on the New York City subway*. 2007. URL: https://www.atlasobscura.com/places/masstransiscope.

[58] Kai-Wen Hsiao, Jia-Bin Huang, and Hung-Kuo Chu. "Multi-view Wire Art". In: *ACM Trans. Graph.* (2018).

[59] Joey Huang et al. "Deepening children's STEM learning through making and creative writing". In: *International Journal of Child-Computer Interaction* 40 (2024).

[60] Yijiang Huang et al. "FrameFab: robotic fabrication of frame shapes". In: *ACM Trans. Graph.* (2016).

[61] Walt Disney Imagineering. *Great Moments with Mr. Lincoln*. 1964.

[62] Walt Disney Imagineering. *Pirates of the Caribbean*. 1967.

[63] Phidgets Inc. *Phidgets*. Accessed on January 5, 2025. URL: https://www.phidgets.com//.

[64] Karen E Jachimowicz and Ronald S Gold. "Stereoscopic (3D) projection display using polarized color multiplexing". In: *Optical Engineering* 29.8 (1990), pp. 838–842.

[65] Alec Jacobson et al. *gptoolbox: Geometry Processing Toolbox*. http://github.com/alecjacobson/gptoolbox. 2018.

[66] Alec Jacobson. "RodSteward: A Design-to-Assembly System for Fabrication using 3D-Printed Joints and Precision-Cut Rods". In: *Comput. Graph. Forum* (2019).

[67] Alec Jacobson, Daniele Panozzo, et al. *libigl: A simple C++ geometry processing library*. http://libigl.github.io/libigl/. 2020.

[68] Henry Jenkins. *Confronting the challenges of participatory culture: Media education for the 21st century*. The MIT press, 2009.

[69] Caigui Jiang et al. "Design and volume optimization of space structures". In: *ACM Trans. Graph.* (2017).

[70] Andrew Jones et al. "Rendering for an interactive 360 light field display". In: *ACM SIGGRAPH 2007 papers*. 2007, 40–es.

[71] Max Longhurst Joshua Boling and Kimberly Lott. "Watersheds, Communities, and Collaboration". In: *Science and Children* 59.4 (2022), pp. 21–25.

[72] Majeed Kazemitabaar et al. "Makerwear: A tangible approach to interactive wearable creation for children". In: *Proceedings of the 2017 chi conference on human factors in computing systems*. 2017, pp. 133–145.

[73] M Khine and Shaljan Areepattamannil. "Steam education". In: *Springer* 10.978-3 (2019), pp. 15–16.

[74] Bogyeong Kim, Jaehoon Pyun, and Woohun Lee. "Enhancing Storytelling Experience with Story-Aware Interactive Puppet". In: *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. 2018, pp. 1–6.

[75] Gou Koutaki. "Eigen Zoetrope". In: *ACM SIGGRAPH 2019 Emerging Technologies*. SIGGRAPH '19. Los Angeles, California, 2019.

[76] Robert Kovacs et al. "Trusscillator: A System for Fabricating Human-Scale Human-Powered Oscillating Devices". In: *The 34th Annual ACM Symposium on User Interface Software and Technology*. 2021, pp. 1074–1088.

[77] Robert Kovacs et al. "TrussFab: Fabricating Sturdy Large-Scale Structures on Desktop 3D Printers". In: *Proc. CHI*. 2017, pp. 2606–2616.

[78] Gregory Kramida. "Resolving the vergence-accommodation conflict in head-mounted displays". In: *IEEE transactions on visualization and computer graphics* 22.7 (2015), pp. 1912–1931.

[79] Sarah Kushner, Paul H. Dietz, and Alec Jacobson. "Interactive 3D Zoetrope with a Strobing Flashlight". In: *Adjunct Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. UIST '22 Adjunct. Bend, OR, USA: Association for Computing Machinery, 2022.

[80] Sarah Kushner et al. "Levitating Rigid Objects with Hidden Rods and Wires". In: *Computer Graphics Forum* (2021).

[81] Sarah Kushner et al. "Papertronic Puppets: Teaching STEM and Storytelling Through Creative Construction". In: *2024 IEEE Frontiers in Education Conference (FIE)*. IEEE. 2024, pp. 1–9.

[82] CREATE Lab. *About Arts & Bots*. Accessed on May 16, 2024. URL: http://www.artsandbots.com/about.html.

[83] Hui Liang et al. "Puppet Narrator: Utilizing Motion Sensing Technology in Storytelling for Young Children". In: *2015 7th International Conference on Games and Virtual Worlds for Serious Applications (VS-Games)*. 2015, pp. 1–8.

[84] Christine Liao. "From Interdisciplinary to Transdisciplinary: An Arts-Integrated Approach to STEAM Education". In: *Art Education* 69.6 (2016), pp. 44–49.

[85] William E. Lincoln. *Toy*. Assignor: Milton Bradley Company. 1876.

[86] Jikai Liu et al. "Current and Future Trends in Topology Optimization for Additive Manufacturing". In: *Struct. Multidiscip. Optim.* (2018).

[87] Liane Makatura, Catherine Most, and Gloria Li. *Balancing Act: An Interactive Tool for Fabricating Calder-Style Hanging Mobiles*. Tech. rep. 2016.

[88] Natalia D Mankowska et al. "Critical flicker fusion frequency: a narrative review". In: *Medicina* 57.10 (2021), p. 1096.

[89] E.J. Marey. *Physiologie du mouvement: vol des oiseaux*. 1890.

[90] John P McIntire, Paul R Havig, and Eric E Geiselman. "What is 3D good for? A review of human performance on stereoscopic 3D displays". In: *Head-and Helmet-Mounted Displays XVII; and Display Technologies and Applications for Defense, Security, and Avionics VI* 8383 (2012), pp. 280–292.

[91] Marshall McLuhan. "The medium is the message". In: *Communication theory*. Routledge, 2017, pp. 390–402.

[92] Vittorio Megaro et al. "Designing cable-driven actuation networks for kinematic chains and trees". In: *Proc. SCA*. 2017.

[93] Sam Mejias et al. "The trouble with STEAM and why we use it anyway". In: *Science Education* 105.2 (2021), pp. 209–231.

[94] Ben Mildenhall et al. "Nerf: Representing scenes as neural radiance fields for view synthesis". In: *Communications of the ACM* 65.1 (2021), pp. 99–106.

[95] M.B. Miles, A.M. Huberman, and J. Saldana. *Qualitative Data Analysis*. SAGE Publications, 2014.

[96] Xiongkuo Min et al. "Sound influences visual attention discriminately in videos". In: *2014 Sixth International Workshop on Quality of Multimedia Experience (QoMEX)*. IEEE. 2014, pp. 153–158.

[97] Niloy J. Mitra and Mark Pauly. "Shadow Art". In: *ACM Trans. Graph.* (2009).

[98] Leo Miyashita et al. "ZoeMatrope: A System for Physical Material Design". In: *ACM Trans. Graph.* (2016).

[99] Eadweard Muybrdge. *Horse in Motion*. 1878.

[100] Anna Newley et al. "Animatronic lions, and tigers, and bears, oh my!" In: *Sci Child* 56 (2019), pp. 64–71.

[101] Anna Newley et al. "Engaging elementary and middle school students in robotics through hummingbird kit with Snap! visual programming language". In: *Journal of Learning and Teaching in Digital Age* 1.2 (2016), pp. 20–26.

[102] Lorelli S Nowell et al. "Thematic analysis: Striving to meet the trustworthiness criteria". In: *International journal of qualitative methods* 16.1 (2017).

[103] Yoichi Ochiai et al. "Fairy lights in femtoseconds: Aerial and volumetric graphics rendered by focused femtosecond laser combined with computational holographic fields". In: *ACM Transactions on Graphics (TOG)* 35.2 (2016), pp. 1–14.

[104] Hyunjoo Oh et al. "FoldMecha: Exploratory Design and Engineering of Mechanical Papercraft". In: *Proceedings of the Eleventh International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '17. Association for Computing Machinery, 2017, pp. 131–139.

[105] OpenAI. *Sora: Creating video from text*. Tech. rep. 2024. URL: https://openai.com/sora.

[106] OpenAI et al. *GPT-4 Technical Report*. 2024. arXiv: 2303.08774 [cs.CL]. URL: https://arxiv.org/abs/2303.08774.

[107] Pauli Pedersen. "Topology optimization of three-dimensional trusses". In: *Topology design of structures*. Springer, 1993, pp. 19–30.

[108] Christopher Thomas Peel et al. "An investigation into the construction of an animatronic model." PhD thesis. University of Bradford, 2010.

[109] Eastern State Penitentiary. *Terror Behind the Walls*. 1991.

[110] Jesús Pérez et al. "Design and fabrication of flexible rod meshes". In: *ACM Trans. Graph.* (2015).

[111] Nico Pietroni et al. "Position-based tensegrity design". In: *ACM Trans. Graph.* (2017).

[112] Romain Prévost et al. "Make it stand: balancing shapes for 3D fabrication". In: *ACM Trans. Graph.* (2013).

[113] Alessandro Giuseppe Privitera, Federico Fontana, and Michele Geronazzo. "The Role of Audio in Immersive Storytelling: a Systematic Review in Cultural Heritage". In: *Multimedia Tools and Applications* (2024), pp. 1–39.

[114] *PyGame Python Library*. https://www.pygame.org/. 2024.

[115] Alec Radford et al. "Learning transferable visual models from natural language supervision". In: *International conference on machine learning*. PMLR. 2021, pp. 8748–8763.

[116] Alec Radford et al. "Robust speech recognition via large-scale weak supervision". In: ICML'23. JMLR.org, 2023.

[117] Ismo Rakkolainen et al. "The interactive fogscreen". In: *ACM SIGGRAPH 2005 Emerging technologies*. 2005, 8–es.

[118] Lukas Rambold et al. "AirTied: Automatic Personal Fabrication of Truss Structures". In: *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 2023, pp. 1–10.

[119] Mitchel Resnick and Eric Rosenbaum. "Designing for tinkerability". In: *Design, make, play*. Routledge, 2013, pp. 163–181.

[120] Jessica Roe et al. *Woo! Jr. kids activities: Children's publishing*. Sept. 2022. URL: https://woojr.com/.

[121] Patrick Ryan. "The storyteller in context: Storyteller identity and storytelling experience". In: *Storytelling, Self, Society* 4.2 (2008), pp. 64–87.

[122] Jesse Schell. *The Art of Game Design: A book of lenses*. CRC press, 2008.

[123] Christian Schüller, Daniele Panozzo, and Olga Sorkine-Hornung. "Appearance-mimicking Surfaces". In: *ACM Trans. Graph.* (2014).

[124] Yuliy Schwartzburg et al. "High-contrast Computational Caustic Design". In: *ACM Trans. Graph.* (2014).

[125] Jamie Simpson Steele. "Becoming creative practitioners: Elementary teachers tackle artful approaches to writing instruction". In: *Teaching Education* 27.1 (2016), pp. 72–87.

[126] Geertje Slingerland et al. "The power of stories: A framework to orchestrate reflection in urban storytelling to form stronger communities". In: *Community Development* 54.1 (2023), pp. 18–37.

[127] Lanny Smoot et al. "An Interactive Zoetrope for the Animation of Solid Figurines and Holographic Projections". In: *ACM SIGGRAPH 2010 Emerging Technologies*. SIGGRAPH '10. Los Angeles, California, 2010.

[128] Kaisa Snellman, Jennifer M Silva, and Robert D Putnam. "Inequity outside the Classroom: Growing Class Differences in Participation in Extracurricular Activities." In: *Voices in urban education* 40 (2015), pp. 7–14.

[129] Ondrej Stava et al. "Stress relief: improving structural strength of 3D printable objects". In: *ACM Trans. Graph.* (2012).

[130] *StoryJumper*. 2024. URL: https://www.storyjumper.com/.

[131] Alan Sullivan. "DepthCube solid-state 3D volumetric display". In: *Stereoscopic displays and virtual reality systems XI*. Vol. 5291. SPIE. 2004, pp. 279–284.

[132] Satu Tenhovirta et al. "Cross-age peer tutoring in a technology-enhanced STEAM project at a lower secondary school". In: *International Journal of Technology and Design Education* 32.3 (2022), pp. 1701–1723.

[133] *Vocaroo - The premier voice recording service*. 2024. URL: https://vocaroo.com/.

[134] Etienne Vouga et al. "Design of self-supporting surfaces". In: *ACM Trans. Graph.* (2012).

[135] Joseph Wachelder. "Toys as Mediators". In: *Icon* (2007), pp. 135–169.

[136] Ingo Wald et al. "Embree: a kernel framework for efficient CPU ray tracing". In: *ACM Trans. Graph.* (2014).

[137] "Walt Disney Audio-Animatronics Timeline". In: *IEEE Potentials* 38.5 (2019), pp. 24–25.

[138] Weiming Wang et al. "Cost-effective printing of 3D objects with skin-frame structures". In: *ACM Trans. Graph.* (2013).

[139] Gordon Wetzstein et al. "Tensor displays: compressive light field synthesis using multilayer displays with directional backlighting". In: (2012).

[140] Frederik Agnar Widjaja. "UNDERSTANDING FLASH FICTION BY THE SIX-WORD STORY: "FOR SALE: BABY SHOES, NEVER WORN""". In: *DIALEKTIKA: JURNAL BAHASA, SASTRA DAN BUDAYA* 7.1 (2020), pp. 56–66.

[141] Andrew J Woods. "Crosstalk in stereoscopic displays: a review". In: *Journal of Electronic Imaging* 21.4 (2012), pp. 040902–040902.

[142] Walt Disney World. URL: https://disneyworld.disney.go.com/en_CA/attractions/magic-kingdom/its-a-small-world/.

[143] Jun Wu, Christian Dick, and Rüdiger Westermann. "A System for High-Resolution Topology Optimization". In: *IEEE Trans. Vis. Comput. Graph.* (2016).

[144] Andrew Wyeth. *Christina's World*. 1948.

[145] Recep Yılmaz and Fatih Mehmet Ciğerci. "A brief history of storytelling: From primitive dance to digital narration". In: *Handbook of research on transmedia storytelling and narrative strategies*. IGI Global, 2019, pp. 1–14.

[146] Tomohiro Yokota and Tomoko Hashida. "Magic Zoetrope: Representation of Animation by Multi-Layer 3D Zoetrope with a Semitransparent Mirror". In: *SIGGRAPH Asia 2018 Emerging Technologies*. SA '18. Tokyo, Japan, 2018.

[147] Cem Yuksel. "Sample Elimination for Generating Poisson Disk Sample Sets". In: *Proc. Eurographics* (2015).

[148] Tomás Zegard and Glaucio H Paulino. "GRAND3: Ground structure based topology optimization for arbitrary 3D domains using MATLAB". In: *Structural and Multidisciplinary Optimization* (2015).

[149] Tomás Zegard et al. "Advancing building engineering through structural and topology optimization". In: *Structural and Multidisciplinary Optimization* (2020).

[150] Xiaoting Zhang et al. "Perceptual models of preference in 3D printing direction". In: *ACM Trans. Graph.* (2015).

[151] Qingnan Zhou, Julian Panetta, and Denis Zorin. "Worst-case structural analysis". In: *ACM Trans. Graph.* (2013).

[152] Kening Zhu and Shengdong Zhao. "AutoGami: a low-cost rapid prototyping toolkit for automated movable paper craft". In: CHI '13. Association for Computing Machinery, 2013, pp. 661–670.

[153] Aleksandar Zivanovic. "The development of a cybernetic sculptor: Edward Ihnatowicz and the Senster". In: *Proceedings of the 5th Conference on Creativity & Cognition*. 2005, pp. 102–108.