

Grasping Information and Collaborating through  
Shape Displays

by

Daniel Leithinger

Submitted to the Program in Media Arts and Sciences,  
School of Architecture and Planning  
in partial fulfillment of the requirements for the degree of  
Doctor of Philosophy in Media Arts and Sciences


at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2015

© Massachusetts Institute of Technology 2015. All rights reserved.

ARCHIVES



**Signature redacted**

Author .....

Program in Media Arts and Sciences  
August 17, 2015


**Signature redacted**

Certified by .....

Prof. Hiroshi Ishii  
Jerome B. Wiesner Professor of Media Arts and Sciences  
Program in Media Arts and Sciences  
Thesis Supervisor

**Signature redacted**

Accepted by .....



Prof. Pattie Maes  
Academic Head  
Program in Media Arts and Sciences



# Grasping Information and Collaborating through Shape Displays

by

Daniel Leithinger

Submitted to the Program in Media Arts and Sciences,  
School of Architecture and Planning  
on August 17, 2015, in partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy in Media Arts and Sciences

## Abstract

The vision to interact with computers through our whole body - to not only visually perceive information, but to engage with it through multiple senses has inspired human computer interaction (HCI) research for decades. Shape displays address this challenge by rendering dynamic physical shapes through computer controlled, actuated surfaces that users can view from different angles and touch with their hands to experience digital models, express their ideas and collaborate with each other. Similar to kinetic sculptures, shape displays do not just occupy, rather they redefine the physical space around them. By dynamically transforming their surface geometry, they directly push against hands and objects, yet they also form a perceptual connection with the users gestures and body movements at a distance. Based on this principle of spatial continuity, this thesis introduces a set of interaction techniques that move between touching the interface surface, to interacting with tangible objects on top, and to engaging through gestures in relation to it. These techniques are implemented on custom-built shape display systems that integrate physical rendering, synchronized visual display, shape sensing, and spatial tracking. On top of this hardware platform, applications for computer-aided design, urban planning, and volumetric data exploration allow users to manipulate data at different scales and modalities. To support remote collaboration, shared telepresence workspaces capture and remotely render the physical shapes of people and objects. Users can modify shared models, and handle remote objects, while augmenting their capabilities through altered remote body representations. The insights gained from building these prototype workspaces and from gathering user feedback point towards a future in which computationally transforming materials will enable new types of bodily, spatial interaction with computers.

Thesis Supervisor: Prof. Hiroshi Ishii

Title: Jerome B. Wiesner Professor of Media Arts and Sciences

Program in Media Arts and Sciences



**Grasping Information and Collaborating through Shape Displays**

by

Daniel Leithinger

**Signature redacted**

Thesis Reader .....

.....

V. Michael Bove Jr.

Principal Research Scientist

MIT Media Lab




**Grasping Information and Collaborating through Shape Displays**

by

Daniel Leithinger

  
**Signature redacted**

Thesis Reader .....



.....

Ivan Poupyrev

Technical Program Lead

Google



# Acknowledgments

Oh where to begin!

My advisor Prof. Hiroshi Ishii is a force of nature who has been an inspiration through his energy and enthusiasm for the last seven years at the Media Lab. I would also like to thank my thesis committee, Michael Bove and Ivan Poupyrev, who have provided me with insights and encouragement. My advisor during my studies in Hagenberg, Prof. Michael Haller, has set me on this path of HCI research and helped me to find a way that has eventually brought me to where I am now.

I have had a chance to work with many amazing people at the Media Lab over the last few years: Sean Follmer was my closest collaborator and a great friend, I hope we have time to discuss new ideas over a bike ride again soon. The Tangible Media Group attracts a special kind of people, and I am glad that I had a chance to be part of this enthusiastic, crazy and lovely crowd. Jamie Zigelbaum, JB Labrune, Amanda Parkes, Keywon Chung, Adam Kumpf, Jinha Lee, Xiao Xiao, Lining Yao, Jifei Ou, Viirj Kan, Ken Nakagaki, Luke Vink, Leonardo Bonanni, Philipp Schoessler, Felix Heibeck, Basheer Tome, Sheng Kai Tang, Samuel Luescher, Austin Lee, Tony DeVincenzi, Cati Vaucelle, David Lakatos, James Patten, and Alex Olwal.

I'd also like to thank Linda Peterson and Keira Horowitz, who have been patient and supportive in navigating the academic maze, especially during the last few weeks. Thanks to Tom Lutz and John DiFrancesco for teaching me many new shop tools, Kevin Davis and Cornelle King for their great support in getting anything around the lab done, and our heroic group administrators Mary Tran Niskala, Lisa Lieberon, Sarah Hunter, Natalia Villegas and Amber Franey.

I don't know how to even start to thank my parents and siblings for supporting me on my travels around the world and occasionally joining me along the way. And at the center of it all, my wife Ateya, my wonderful lover, companion, and best friend.



# Contents

<b>1</b>	<b>Introduction</b>	<b>23</b>
1.1	Approach . . . . .	26
1.2	Contributions . . . . .	26
1.3	Motivation and Project History . . . . .	27
1.3.1	Statement on Multiple Authorship and Prior Publications . . . . .	30
1.4	Thesis Overview . . . . .	31
<b>2</b>	<b>Background and Related Work</b>	<b>33</b>
2.1	Embodied Cognition . . . . .	34
2.2	Kinetic and Cybernetic Art . . . . .	35
2.3	Virtual Reality (VR) . . . . .	41
2.4	Volumetric Displays . . . . .	42
2.5	Haptic Interfaces . . . . .	43
2.6	Tangible User Interfaces . . . . .	46
2.7	Shape-Changing User Interfaces . . . . .	47
<b>3</b>	<b>A Primer on Shape Displays</b>	<b>51</b>
3.1	Chronological Development . . . . .	52
3.2	Purpose and Application Domains . . . . .	54
3.2.1	Art . . . . .	54
3.2.2	Data Physicalization . . . . .	56
3.2.3	Remote Communication . . . . .	59
3.2.4	CAD Tool . . . . .	59

3.2.5	Accessibility . . . . .	60
3.2.6	Dynamic Physical Controls . . . . .	60
3.2.7	Material Handling . . . . .	61
3.2.8	Other Applications . . . . .	61
3.2.9	Analysis of Application Domains . . . . .	62
3.3	Parameter Space for Shape Rendering . . . . .	65
3.3.1	Dimensions and Form Factor . . . . .	65
3.3.2	Resolution . . . . .	66
3.3.3	Actuation Speed . . . . .	67
3.3.4	Vertical Range and Resolution . . . . .	67
3.3.5	Haptic Feedback and Actuation Force . . . . .	67
3.3.6	Sound . . . . .	68
3.3.7	Interface Surface . . . . .	68
3.3.8	Graphics . . . . .	69
3.4	Representing Content . . . . .	70
<b>4</b>	<b>Interaction</b>	<b>75</b>
4.1	Input Modalities . . . . .	77
4.1.1	Touch . . . . .	77
4.1.2	Deformation . . . . .	78
4.1.3	External Controllers . . . . .	79
4.1.4	Mid-Air Gestures . . . . .	79
4.1.5	Tangible Objects . . . . .	81
4.2	Dynamic Scale and Modality . . . . .	81
4.3	Dynamic Scale with ShapePorts . . . . .	85
4.3.1	Interaction Techniques . . . . .	86
4.3.2	Providing Context with Graphics Displays . . . . .	87
4.3.3	Applications . . . . .	88
4.4	Dynamic Modality with Sublimate . . . . .	90
4.4.1	Related Work in Augmented Reality . . . . .	94

4.4.2	Interactions . . . . .	97
4.4.3	System Setup . . . . .	99
4.4.4	Applications . . . . .	102
4.4.5	Single User Applications . . . . .	103
4.4.6	Multi-User Applications . . . . .	105
<b>5</b>	<b>Remote Collaboration through Physical Telepresence</b>	<b>107</b>
5.1	Related Work in Distributed Collaboration . . . . .	109
5.1.1	Video-Mediated Communication . . . . .	109
5.1.2	Mixed Reality . . . . .	109
5.1.3	Telerobotics . . . . .	110
5.1.4	Tangible Remote Collaboration . . . . .	110
5.1.5	Shape Displays for Remote Collaboration . . . . .	111
5.2	Physical Telepresence . . . . .	111
5.2.1	Example Scenario: Remote 3D design collaboration . . . . .	112
5.2.2	Remote Physical Embodiment . . . . .	113
5.2.3	Representations: Physical and Virtual . . . . .	114
5.3	Physically Manipulating Remote Objects . . . . .	115
5.3.1	Gesture Control . . . . .	117
5.3.2	Interface Elements . . . . .	118
5.3.3	Tangibles and Physical Tools . . . . .	118
5.4	Computer-Mediated Shape I/O . . . . .	119
5.4.1	Transformation of Physical Form: Bending Body Limits . . . . .	120
5.4.2	Altered Representation: Becoming Something Else . . . . .	120
5.4.3	Filtering: Adding and Removing Motion . . . . .	122
5.4.4	Motion Dynamics: Warping Time . . . . .	122
5.5	Prototype Applications . . . . .	122
5.5.1	Telepresence Workspace . . . . .	122
5.5.2	Collaborative 3D Modeling . . . . .	123
5.5.3	Shared Presence: Connected Membranes . . . . .	124

5.6	Technical Implementation . . . . .	125
5.6.1	Shape Output . . . . .	125
5.6.2	Sensing Shape Deformation Input . . . . .	127
5.6.3	Tracking People and Objects with a Depth Camera . . . . .	128
5.6.4	Latency and Manipulating Objects . . . . .	128
5.7	Evaluation . . . . .	129
5.7.1	Apparatus . . . . .	129
5.7.2	Participants . . . . .	130
5.7.3	Task, Conditions . . . . .	130
5.7.4	Procedure . . . . .	131
5.7.5	Study Results and Observations . . . . .	131
<b>6</b>	<b>Implementation</b>	<b>135</b>
6.1	Actuation Platform . . . . .	136
6.1.1	Suitable Actuation Technologies . . . . .	137
6.2	Table Hardware and Setup . . . . .	140
6.2.1	Relief . . . . .	140
6.2.2	inFORM . . . . .	142
6.2.3	TRANSFORM . . . . .	144
6.2.4	inFORM for “Tools” exhibit . . . . .	144
6.2.5	Rendering . . . . .	145
6.2.6	User and Object Tracking . . . . .	145
6.2.7	Future Touch Detection . . . . .	148
<b>7</b>	<b>Discussion</b>	<b>149</b>
7.1	Design Guidelines . . . . .	149
7.1.1	Interaction and Rendering beyond the Surface . . . . .	149
7.1.2	Improvising through Inter-material Interaction . . . . .	150
7.1.3	Animation Principles . . . . .	150
7.2	Tools for Artists and Researchers . . . . .	152
7.2.1	A Platform for Design Fiction . . . . .	152

7.2.2	New Brushes for a Transforming Canvas . . . . .	153
7.3	Technical Improvements . . . . .	159
<b>8</b>	<b>Conclusion</b>	<b>163</b>



# List of Figures

1-1	Compared to the rich interactions with objects and people in the real world, using computers often feels like handling “Pictures Under Glass” [214]. . . . .	23
1-2	Interacting with shape displays through touch and mid-air gestures, and manipulating objects. . . . .	24
1-3	Our first shape display, <i>Relief</i> , was motivated by the earlier <i>SandScape</i> system. . . . .	27
1-4	The <i>Radical Atoms</i> vision as a successor of <i>Tangible Bits</i> [80] . . . . .	29
2-1	Procedures for haptic exploration of object attributes [98] . . . . .	36
2-2	Kinetic art precursors to shape displays. . . . .	38
2-3	VR and Shape Displays placed on the Reality-Virtuality Continuum (graphic adapted from Milgram and Kishino [122] VR photo [130] . . . . .	41
2-4	Multiple users interacting with a volumetric displays [55] . . . . .	43
2-5	Haptic Workstation at EPFL [171] . . . . .	44
2-6	Examples for different technical approaches to shape generation: modular (2-6a), embedded (2-6b) and external (2-6c). . . . .	49
3-1	Interactive creatures on shape displays. . . . .	56
3-2	Examples of Data Physicalization . . . . .	57
3-3	Examples of GIS Applications . . . . .	57
3-4	Parameters for Representation: a) Area b) Pin Spacing c) Pin Diameter d) Vertical Range e) Speed f) Haptic Feedback h) Sound . . . . .	64
3-5	Examples of different interface surfaces on <i>Relief</i> . . . . .	68

3-6	Augmenting the interface surface with graphics: a) top projection [113] b) rear projection [180] c) embedded [152] d) projection mapped onto sides . . . . .	69
3-7	Additional graphics displays: a) head-mounted b) hand-held c) see through d) projection mapped e) vertical display . . . . .	70
3-8	Rendering a 3D model on a shape display usually results in geometry with missing overhangs, and clipping and interpolation errors. . . . .	71
4-1	Interactions with bare hands through touch and deformation. . . . .	75
4-2	Interactions beyond direct touch: gestures and tangible objects. . . . .	76
4-3	Limitations of direct touch interactions with shape display. . . . .	78
4-4	Mid-air gestures for selecting and translating content. . . . .	80
4-5	Mid-air gestures for scaling and rotating content. . . . .	80
4-6	Multiple viewports and rendering styles in modern CAD software. . . . .	82
4-7	Building a physical architectural model. (Image courtesy by Denise Coté.) . . . . .	82
4-8	Interaction Models of GUI and TUI (adapted from [207]) . . . . .	83
4-9	Concept sketch to show the viewport and the modality of the dynamic tangible representation. . . . .	84
4-10	Possible viewport transformations. . . . .	87
4-11	Combining a ShapePort with a surrounding tabletop projection and wall displays for context. . . . .	88
4-12	CityScape setup: shape display with a wall display rendering a contex- tual view. . . . .	89
4-13	Various data sets rendered in CityScape. . . . .	89
4-14	Concept sketch of a user interacting with a partly physical model on the <i>Sublimate</i> system. . . . .	91
4-15	The transitions between <i>spatial 3D graphics</i> and <i>shape output</i> are in- spired by phase transitions. . . . .	92
4-16	Sublimate setups. . . . .	100

4-17	NURBS Surface Modeling. The user can switch between rendering a) surface physically and control points virtually or b) control points physically and the surface virtually. . . . .	102
4-18	Volumetric Medical Data Viewing. Users can modify cross-sections through the volume by physically deforming the object shape and switch between defined cross-sections through sublimation. . . . .	104
4-19	Virtual Wind Tunnel. The wind flow around a physical model formed by the user is visualized through graphics. . . . .	105
4-20	Multi-user geospatial data exploration. Hand-held tablets augment the shape display by adding layers of data, and extending the active workspace. . . . .	106
5-1	Co-located and remote collaboration through shape displays. . . . .	108
5-2	A user manipulating a sphere through his remote physical embodiment.	112
5-3	Top: Two linked tangible tokens at two locations, with synchronized position. Bottom: A tangible token rendered with the shape display at a remote location. . . . .	114
5-4	Improvising with physical objects: Using a basket to scoop up remote balls. . . . .	116
5-5	Transmitting gesture shapes directly above the surface, or from a sep- arate mid-air zone above the shape output. . . . .	116
5-6	The brush is equipped with a switch to detect when it is pushing against the side of a rendered object. . . . .	118
5-7	Pushing against the side of a shape-rendered objects with the brush tool moves it. . . . .	119
5-8	Scaling a users hand to interact with larger objects. . . . .	120
5-9	Copying and rendering multiple hands simultaneously (left) and switch- ing to a virtual, non-physical, representation (right). . . . .	121
5-10	Replacing hands with a hook to reach or ramps to slide objects. . . . .	121
5-11	The claw tool open and closed to enclose and move an object. . . . .	121

5-12	Remote assistance: A remote expert helps a local user by supporting the workpiece with one hand and pointing with the other. . . . .	123
5-13	Collaborative 3D modeling of two linked landscape models. . . . .	124
5-14	Connected Membrane: When a user pushes into the membrane on one site, the inverse shape appears at the other. . . . .	125
5-15	Side view of the TRANSFORM system with a single 12x2 actuation module. . . . .	126
5-16	System diagram of the telemanipulation system using shape capture and remote shape display. . . . .	126
5-17	System diagram of connected workspaces with bidirectional shape capture and shape rendering. . . . .	127
5-18	Setup of the telemanipulation study. Left: The participant’s gestures are captured and two screens provide visual feedback. Right: Shape output moves the target ball. . . . .	129
5-19	Participant’s view of the remote scene, with a top down 1:1 view on the horizontal screen, and a perspective view on the vertical. . . . .	130
5-20	Averages of users’ ratings of task difficulty for 4 different input conditions in each task. (7-point likert scale, 7 is best/easiest). Error Bars are standard error of the mean (SEM). . . . .	132
6-1	Hardware setups for inForm and Transform shape displays. . . . .	135
6-2	RSA0N11M Motor-driven master type slide potentiometer by ALPS [2].	140
6-3	The inFORM system actuates and detects shape change with 900 mechanical actuators, while user interaction and objects are tracked with an overhead depth camera. A projector provides visual feedback. . . .	143
6-4	The inFORM setup created for the “Tools” exhibit. . . . .	145
6-5	Basic pipeline for rendering shapes and graphics and tracking user input with a depth camera. . . . .	146
6-6	Tracking pipeline to detect the users gestures and touch. . . . .	147
6-7	Tracking pipeline to detect object positions. . . . .	148

7-1	Still image of the video prototype <i>TRANSFORM as Dynamic and Adaptive Furniture</i> [216]. . . . .	153
7-2	TRANSFORM timeline GUI for triggering and blending animations.	155
7-3	<i>3ds Max</i> networking plugin to play back animations on <i>TRANSFORM</i> .	156
7-4	An animator deforming a fabric surface to inform shape change [38]. .	157
7-5	A wave animation is driven by the user's mid-air gestures. . . . .	158



# Chapter 1

## Introduction

Computers provide us with access to an unprecedented and ever-increasing amount of information, with algorithms at our fingertips we can investigate data, create new models and exchange information with each other like never before. But while computers have progressively become more mobile, connected and capable, their physical interfaces are still dominated by rigid, flat, displays that shape how we perceive and manipulate information. Bret Victor has accurately described the resulting interaction as handling “Pictures Under Glass”, criticizing how these interfaces compromise the expressive capabilities of the human hands and body [214]. There is an emerging trend in computing to move beyond the traditional desktop metaphor with new interfaces and form factors. Mobile devices and touch screens liberate and enrich




(a) Child’s hands in clay [166]



(b) Interacting with a tablet [29]

Figure 1-1: Compared to the rich interactions with objects and people in the real world, using computers often feels like handling “Pictures Under Glass” [214].



(a) Touching the representation of a surface equation. (b) Gestural interaction with a dynamic surface. (c) Remote manipulation of an object.

Figure 1-2: Interacting with shape displays through touch and mid-air gestures, and manipulating objects.

our interactions through multitouch, motion sensors and gestures. Virtual Reality (VR) technology is becoming widely available thanks to advances in motion sensing, display technology and mobile hardware, which has enabled devices like *Oculus Rift* and *Google Cardboard* to start bringing decades of VR research into the consumer market. The spatial interaction in these systems, however, focuses primarily on the visual senses – but what about our rich abilities to perceive, understand, and more importantly, interact and manipulate the world through our hands? Our current relationship to data is best summed up by paraphrasing Malcolm McCullough: “Eyes are in charge, but hands are underemployed” [119]. To challenge this dominance of pixels, human computer interaction (HCI) research has for decades been inspired by the vision to engage with digital information through our body in richer and more complex ways, by engaging the sense of touch. A popular approach to achieve this goal are haptic devices, computer interfaces that render touch sensations.

Most haptic devices render sensations on a few discrete, limited points of contact; in contrast, *shape displays* generate physical shapes through computer controlled, actuated surfaces that users view from various angles and touch with their hands. By interacting with these dynamic shapes, multiple users can experience digital models, express ideas, and collaborate with each other in tangible ways. But while these interfaces hold great promise for human computer interaction, their application domains are not yet well defined and few interaction techniques for them have been developed so far. At this point, there are two fundamental questions for researchers to address:

First, why should we build shape displays? Second, how can we build them? The second question pertaining to the engineering challenge to build shape displays that are high-resolution, robust, and have many degrees of freedom has been addressed by related work before and is an active and evolving research area. But the more fundamental question remains open: Why should we build shape displays, and what advantages do they offer over simpler approaches for haptic rendering and tangible interaction? To answer this question, this thesis explores how shape displays enable rich spatial interaction that combines tangible interaction with dynamic shapes, objects and mid-air gestures. By generating dynamic physical forms, shape displays do not just occupy, but rather they redefine the physical space around them, much like kinetic sculptures. Through their transforming surface geometry, they directly push against hands and objects, but also form a perceptual connection with the users gestures and body movements at a distance. I define this transition from physical to perceptual connection as *spatial continuity*. Based on this principle of spatial continuity, this project defines a set of interaction techniques that move between touching the interface surface, to manipulating tangible objects on top, and to interacting through gestures in the space above it (see Figure 1-2).

Based on an analysis of application domains for shape displays, we define a set of interactions for manipulating data at different *scales* and *modalities*. We propose interactions to physically zoom in onto different levels of detail of a data model, pan across it, or rotate the physical representation. These transformations help to overcome some limitations of current shape display hardware, like size, range, and degrees of freedom, while placing data within reach of the user's physical abilities. In addition to dynamic scale, we propose dynamic modality, through a system that extends a shape display with spatially co-located 3D graphics. Current shape displays trade the advantage of real materials for the flexibility and realism of high-resolution graphics of virtual reality and augmented reality interfaces. Combining these two modalities provides a rich computing experience neither one of them could enable on its own. Particularly interesting are interactions based on the ability to rapidly switch between rendering objects as a solid shape in one moment and as an intangible

floating graphic in the next, or as a hybrid of both modalities.

To demonstrate how these interaction techniques can be utilized across various application domains, we apply them to a range of applications for computer-aided design, urban planning, and volumetric data exploration. To support remote collaboration, we introduce a new approach to telepresence through shared workspaces with the ability to capture and remotely render the physical shapes of people and objects. Based on the concept of shape transmission, we propose techniques to manipulate remote objects and linked physical renderings of digital content. We investigate, how the representation of users' body parts can be altered to amplify their capabilities for teleoperation and study how users spatially interact through a prototype physical telepresence workspace.


## 1.1 Approach

This thesis presents three shape display hardware prototypes: *Relief*, *inFORM* and *Transform*. These systems integrate physical rendering, synchronized visual display, shape sensing, and spatial tracking of objects and gestures. Based on this platform, my collaborators and I have built a series of tangible applications for computer aided design, volumetric data exploration and remote collaboration. The insights gained from building these systems and testing them with end users have resulted in a set of tools for designing shape change and in guidelines for application development to help bring these interfaces out of research labs and into the hands of everyday computer users.

## 1.2 Contributions

This thesis presents the following contributions:

- An analysis of the applications and design space of content representation through shape displays.
- A framework for interaction techniques for shape displays using touch, gestures,



(a) The *SandScape* system for landscape design [78]. (b) A user touching a landscape model on the *Relief* shape display.

Figure 1-3: Our first shape display, *Relief*, was motivated by the earlier *SandScape* system.

and tangible objects.

- A technical description of novel shape display hardware: *Relief*, *inFORM*, and *TRANSFORM*.
- A system to augment shape output with spatial graphics: *Sublimate*.
- A description of object manipulation through shape deformation.
- A system for remote collaboration, which transmits the physical shape of people, objects and shared content.
- A study of interaction techniques for the manipulation of remote physical objects through shape displays.
- A set of tools for designing shape change.
- User interface guidelines for the application development of shape displays.

### 1.3 Motivation and Project History

The research presented in this thesis started with the *Relief* project in 2009 (see Figure 1-3b), which was motivated by a number of earlier projects in the Tangible Media Group. A major source of inspiration was the *Illuminating Clay* project [147] and the subsequent *SandScape* project by Piper et al. (see Figure 1-3a) [78], which to me are still amongst the most engaging examples of a Tangible User Interface (TUI) to date.

These interfaces for computer-aided design (CAD) allow users to create a landscape model by forming real, physical material with their hands, and by improvising with tools and materials during this creative process. The system continuously scans this physical landscape model and updates a corresponding digital model. It then projects graphical feedback of computed model parameters such as slope, shadow casting, and local drain direction back onto the physical model. This tight coupling between rich material input and projected graphical output successfully blurs the boundary between the physical and the digital world: designers feel like they are interacting with a physical model that has digital capabilities. However, the use of inert, physical materials for interaction imposes several functional limitations compared to modern CAD software. The system can only sense user input through the landscape model, but is unable to modify the physical model shape; therefore operations like loading a model from a file, rendering models that change over time, and changing global model parameters like scale result in a discrepancy between the digital and the physical model.

The idea of using shape displays to overcome such limitations through physical shape rendering has been previously explored by other research groups (as outlined in Section 3.1) and within the Tangible Media Group: The *PINS* project by Kaanta investigated the feasibility of pin-based shape rendering [91], while AR-Jig by Anabuki et al. combined an array of 12 linear actuators with Augmented Reality graphics for a haptic CAD interface [3]. *Relief* was initiated after finding old prototype actuators from the *PINS* project in our group space, which motivated the idea of building a motorized next generation of *Illuminating Clay*. This influence of *Illuminating Clay* and *SandScape* is still evident in some of the design choices for the project, like the tabletop form factor, overall dimensions, and ceiling-mounted projector.

But beyond CAD, *Relief* was conceived as a more general purpose research platform: at the time the project started, the Tangible Media Group hosted numerous internal discussions about the future of *Tangible User Interfaces*, which eventually led to the research vision of *Radical Atoms*, defined by my advisor Hiroshi Ishii as a successor of the *Tangible Bits* vision (see Figure 1-4) [80]. *Radical Atoms* is a vision




Figure 1-4: The *Radical Atoms* vision as a successor of *Tangible Bits* [80]

of human interaction with a hypothetical future material that is computationally reconfigurable. As this ideal material does not yet exist, our research goal was to use a shape display as a primitive first platform to propose and explore the interactions with such a future material. This platform would allow us to build functioning prototype systems that could change shape just well enough to explore a future in which computational objects and materials in our everyday environment would be able to drastically transform their shape.

Another influence came from our groups' research in mid-air gestural interfaces based on the *G-Speak* system [229], which spurred experiments to utilize gestures as an additional interaction modality for shape change. In Fall 2009, my colleague Sean Follmer and I collaborated on a stop-motion video prototype called *Contour*, in which we explored interaction with a higher-resolution shape display through mid-air gestures and tangible objects for CAD. A simplified version of these ideas was implemented with the *Recompose* system by my colleagues Matt Blackshaw, David Lakatos, and Tony Devincenzi in Fall 2010, based on the newly available *Kinect* sensor and the *Relief* shape display hardware [12]. Combining *Relief* with spatial 3D graphics and the *G-Speak* gesture tracking system, my colleagues Sean Follmer, Samuel Luescher, Alex Olwal, Akimitsu Hogge, Jinha Lee, and I developed the *Sublimate* system in 2012. As the physical resolution of *Relief* was too low to implement the interactions

envisioned in *Contour*, Sean Follmer and I started to design the *inFORM* shape display in 2012. On this platform we investigated *Dynamic Affordances and Constraints* together with Alex Olwal and Akimitsu Hogge in 2013 [37], and *Physical Telepresence* together with Alex Olwal in 2014 [111]. In the meanwhile, our shape display research unexpectedly reached a large non-academic audience through a video of the *inFORM* system, which between Fall 2013 and Summer 2015 has been viewed about 9 million times across *YouTube* and *Vimeo*, after which the project was covered by various news outlets. Another chance to present shape displays to a non-academic audience was the *TRANSFORM* project, shown at the *Milan Design Week 2014* as part of the *Lexus Design Amazing* exhibit [81]. A collaboration between Hiroshi Ishii, myself, Sean Follmer, Amit Zoran, Jared Counts, and Philipp Schoessler, *TRANSFORM* presented shape change in the context of dynamic furniture and as a storytelling platform. This idea of shape-changing surfaces in our everyday lives was further developed as a video prototype on the *TRANSFORM* system by my colleagues Luke Vink, Viirj Kan, and Ken Nakagaki [216]. We also exhibited an updated version of *inFORM* at the Cooper Hewitt National Design Museum in New York as part of the exhibit: “Tools, Extending Our Reach” from December 2014 through May 2015. All of the above projects were completed under the guidance of my advisor Hiroshi Ishii.

### **1.3.1 Statement on Multiple Authorship and Prior Publications**

The research projects presented in this dissertation were conceived and implemented with my colleagues at the Tangible Media Group, led by Hiroshi Ishii. Sean Follmer was a close collaborator, with whom I ideated, developed and tested most of the research presented in this dissertation. Alex Olwal provided research guidance and co-authored many of the publications listed below. To reflect the collaborative nature of this work, I will use the plural form “we”, when referring to our projects. Our research ideas have previously published in the following conference publications, which form the basis for this dissertation:

- “Relief: A Scalable Actuated Shape Display”, presented at TEI 2010 [113]
- “Recompose: Direct and Gestural Interaction with an Actuated Surface”, presented at CHI 2011 [12]
- “Direct and Gestural Interaction with Relief: a 2.5D Shape Display”, presented at UIST 2011 [114]
- “Sublimate: State-changing Virtual and Physical Rendering to Augment Interaction with Shape Displays”, presented at CHI 2013 [112]
- “inFORM: Dynamic Physical Affordances and Constraints Through Shape and Object Actuation”, presented at UIST 2013 [37]
- “Physical Telepresence: Shape Capture and Display for Embodied, Computer-mediated Remote Collaboration”, presented at UIST 2014 [111]
- “TRANSFORM: Embodiment of Radical Atoms at Milano Design Week”, presented at CHI 2015 [81]
- “TRANSFORM As Adaptive and Dynamic Furniture”, at CHI 2015 [216]

Apart from the research presented in this dissertation, Sean Follmer utilized the *inFORM* shape display system for his dissertation research “Dynamic Physical Affordances for Shape-Changing and Deformable User Interfaces” [38]. *inFORM* was also used as the hardware platform for Philipp Schoessler’s thesis “Shape Synthesis: Physical Object Augmentation and Actuation for Display and Interaction on Shape Changing Interfaces” [175], parts of which will be presented at UIST 2015 as the paper: “Kinetic Blocks: Actuated Constructive Assembly for Interaction and Display”.

## 1.4 Thesis Overview

This thesis consists of an analysis of the background and related work in Chapter 2 and a primer on shape displays in Chapter 3. Chapter 4 proposes interaction techniques and discusses their applications to a series of prototype systems. Finally, Chapter 5 describes shape display systems for remote collaboration through physical telepresence, and Chapter 7 discusses the lessons learned and presents guidelines and future research directions.



## Chapter 2

# Background and Related Work

When building shape displays, the most fundamental questions that related work can help to answer are why it is important to build them, and where we can draw inspiration from when designing them. This chapter discusses the theoretical foundations of shape displays through *Embodied Cognition*. Embodied cognition explains how humans experience the world around them and make sense of it through their body. This motivates why we want to provide the ability to interact with computers through real shapes. This bodily experience of ideas and systems through physical objects has also been explored in the context of *Kinetic Art*. I trace the rich history of kinetic sculptures to understand how they relate to shape displays and how HCI research can benefit from the different angle that art discourse provides for understanding the human experience of shape-change.

Shape displays are closely related to other approaches in HCI that deviate from the *Windows, Icons, Menus and Pointer (WIMP)* paradigm of current Graphical User Interfaces: these including *Virtual Reality*, *Haptic Interfaces*, *Tangible User Interfaces*, and other types of *Shape-Changing User Interfaces*. Other, more project-specific related work is discussed in subsequent chapters of this dissertation: I provide an in-depth analysis of prior shape displays research in Chapter 3, touching upon related work in *Augmented Reality (AR)* in Section 4.4.1 and in *Remote Collaboration* in Section 5.1.

## 2.1 Embodied Cognition

Humans build a cognitive model of their environment and objects around them by actively exploring the world with their body and through their combined senses. Gibson has defined the concept of *active touch* as the activity of a person scanning objects through the act of *touching* [48], employing both cutaneous and kinesthetic sense. He differentiated it from *passive touch*, when a person is being touched. Through a series of comparative studies, he had demonstrated that test subjects who actively explored a shape with their fingers were able to identify objects better than when the shape was only pressed against their passive hand. Gibson's research on *active touch* was extended by Klatzky et al., who classified different hand configurations that humans employ when probing objects. They arrived at a set of *exploratory procedures* (see Figure 2-1), along with a classification of the types of information about an object that a person acquires through touch exploration [98].

Our research draws from the concept of affordances, which Gibson introduced as "what [an object or environment] offers the animal, what it provides or furnishes, either for good or ill.", and which can be viewed as the set of action potentials for an object [50]. Norman first applied the concept of affordances to design and HCI, focusing on "perceived affordances" that the designer creates to provide interaction clues, or suggestions, to the user [133]. Gaver defined *technology affordances* as "properties of the world that are compatible with and relevant for people's interactions." [45] He emphasized the importance of perceptible affordances, since mismatched or hidden affordances interfere with the interface's legibility, which may confuse the user and result in improper operation. Gaver highlights that these affordances can be perceived visually, tactilely, or aurally. He also expanded on sets of affordances: *nested affordances*, which are grouped spatially, and *sequential affordances*, which are grouped temporally. Sequential affordances are often used in GUI based interaction, where graphical perceived affordances can be rendered quickly and then disappear, but are also highly relevant for the affordances of physical user interfaces that transform over time, like shape displays.

Hartson elaborated on Norman and Gaver's work describing four types of affordances: *cognitive affordance* as a "design feature that helps users in knowing something", *physical affordance* as a "design feature that helps users in doing a physical action in the interface", *sensory affordance* as a "design feature that helps users sense something" and *functional affordance* as a "design feature that helps users accomplish work (i.e., the usefulness of a system function)" [63]. Kaptelinin et al. further split both cognitive and physical affordances in two parts, describing the *handling affordances*, the affordances of the part of a tool that the user interacts with, and the *effector affordances*, the affordances of the tool that it manipulates an object with, and suggest that these two must be tightly coupled [94]. Grounded in the philosophical tradition of phenomenology, Dourish proposes the importance of *Embodied Interaction* for HCI [27]. In a similar fashion, Klemmer et al. discuss the importance of our physical bodies to experience the world, understand it and interact with it, and propose themes to inspire new design approaches that integrate the physical world in which our body exists with the digital world of computers [99].

This research in cognition and its application to HCI provides a theoretical foundation and argument for why shape displays can allow for richer *embodied interaction* with information, providing cognitive benefits to the user. Gibson's work on *active touch* has directly motivated the invention of shape displays that allow users to freely move their hands along the interface surface, affording a significantly enhanced cognitive experience of rendered objects than that afforded by the limited touch points of comparable haptic devices [153]. Shape displays generate dynamic physical, cognitive and functional affordances for content and user interface elements to guide interaction. While the user interacts with these features, have the ability to transform over time to provide sequential affordances as discussed by Gaver.

## 2.2 Kinetic and Cybernetic Art

The influence of kinetic art to shape display research is important in two ways: kinetic sculptures are precursors to their form factor, with sculptures by artists like Gianni



Figure 2-1: Procedures for haptic exploration of object attributes [98]

Colombo and Ward Fleming having explored similar concepts decades before the HCI community adopted them. But kinetic art theory also provides us with a discourse to understand how a kinetic object transforms the space around it, and the effect on the human observer. Kinetic art and shape displays share a common notion in physically expressing ideas and systems in space, through object movement.

After painters and sculptors of the Impressionist generations, such as Manet, Monet, Degas, and Rodin had found new ways of expressing light and motion in their art during the 19th century; the representation of movement became a dominant theme in modern art during the early 20th century. From 1910 onwards, cubists like Picasso and Braque revolutionized the plastic arts by conveying the motion of the observer around an object through multiple simultaneous viewpoints. Futurists like Boccioni and Balla discussed and emulated the abstract idea of movement to develop a visual language on how to express and capture the experience of motion through painting and sculpture, rather than producing a literal perception of movement [151].


Kinetic art emerged from these influences as an independent art form in the 1920s; it depicts movement through the literal three-dimensional motion of objects and light reflections in space. Early experiments include Marcel Duchamp's first ready-made *Bicycle Wheel* in 1913, Naum Gabo's *Standing Wave* in 1920, and Lazlo Moholy-

Nagy's *Lichtrequisit* in 1930. The 1920 *Realist Manifesto* by Gabo and Pevsner proclaimed a new artistic expression through kinetic rhythms [42]: "We renounce the thousand-year-old delusion in art that held the static rhythms as the only elements of the plastic and pictorial arts. We affirm in these arts a new element the kinetic rhythms as the basic forms of our perception of real time." Such ideas were expressed in sculptures like *Standing Wave*, which created a volume in space by spinning a standing rod rapidly.


The 1955 group exhibit "Le Mouvement" at the Gallerie Denise Rene in Paris marked the rise of kinetic art as a popular art form. Besides showing the works of Calder and Duchamp, it included many later well-established artists in the field, like Yaacov Agam, Pol Bury, Robert Jacobson, Jesus-Raphael Soto, and Jean Tinguely, and Victor Vasarely. Jean Tinguely's self destructing "Homage to New York" in 1960 at the MOMA sculpture garden was also the first collaboration with Bill Klüver, an engineer from Bell Labs, and artist Robert Rauschenberg. Klüver and Rauschenberg were later instrumental in the organization of the 1966 events "9 Evenings: Theatre and Engineering" at the Sixty-ninth Regiment Armory in New York, and the founding of the "Experiments in Art and Technology (E.A.T)" organization in 1967, which operated until the 1980s.

Kinetic artworks of the early period either used predefined motion through motors, or physically reacted to touch and air currents. Later artists like Nicolas Schöffer, Wen-Ying Tsai and Hans Haacke explored more complex systems with embedded natural or computational feedback mechanisms. These concepts had been applied to the engineering of machines like steam engines earlier, but research on missile trajectory control initiated the development of computed feedback theory during WW2. The subsequent emergence of the field of *Cybernetics*, defined and popularized by Norbert Wiener as "the scientific study of control and communication in the animal and the machine", resulted in the discovery of how natural systems often follow a similar mechanism of regulation and can be computationally modeled [220].

Nicolas Schöffer, who created kinetic sculptures since 1949, is widely regarded as the father of cybernetic art, but was also on the intellectual forefront of cybernetic



(a) “Strutturazione pulsante” (1983) by Gianni Colombo [106]



(b) “Pinscreen toy” (1985) by Ward Fleming [76]

Figure 2-2: Kinetic art precursors to shape displays.

architecture. Starting with *CYSP 1* (1956), the first cybernetic sculpture utilizing a computer, he developed large-scale cybernetic sculptures like a 52 m high tower for the city of Liege (1961), while envisioning how an emerging form of cybernetic architecture could improve urban life by reacting to the city in real time and for instance reconfigure to prevent riots. This vision of cybernetic architecture went further than using computation only during the design stage - the buildings themselves were to physically react to people and the environment [19].

This concept was of architecture governed by a cybernetic system and reacting to its environment was also explored by the *Architecture Machine* group at MIT, headed by Nicholas Negroponte. The *Seek* project, on display at the 1970 exhibit “Software” at the Jewish Museum in NY, demonstrated this idea through of a 5 ft x 8 ft glass box of computationally placed 2 inch cubes, with gerbils roaming in-between [90]. The animals would move and knock over the blocks, while a computer-controlled robot arm would attempt to sense these physical changes and resolve inconsistencies with its digital model, trying to restore the blocks, but also adapt to the animals.

Shape displays have been influenced by art, and some of the earliest examples of shape displays are kinetic sculptures, with George Rickey mentioning the possibility of a shape changing canvas as a possible form of kinetic art as early as 1963 [162]. Other kinetic artists started to explore non-computerized form factors similar to shape displays early on, with examples like Gianni Colombo’s “Superficie Pulsante”

(1959), a grid of styrofoam blocks transformed by underlying actuators (see Figure 2-2a for a 1983 larger-scale version), or Ward Fleming’s “pin-table display” (1976), first shown at the SF exploratorium and later developed into the popular pinscreen toy [34], shown in Figure 2-2b. Continuing this heritage, modern shape displays still often have a dual identity as computer interface research prototypes and kinetic artworks. In his paper on “Art and Technology in Interface Devices”, Iwata discusses his motivation for presenting haptic research like the *FEELEX I* shape display outside the academic context at art venues like “Ars Electronica” [85], as it allows him to explore unconventional applications, present them to the public and evaluate them with a diverse audience.

Besides taking inspiration from kinetic sculptures, art theorists provides us with a better understanding of the effects of shape change on the user. While the HCI research and discourse has largely focused on their functional advantages as a *haptic interface*, we are still at the beginning of understanding the different ways in which these interfaces can transform how we interact with data spatially. By discussing shape displays as kinetic artworks, we get a more nuanced perspective on how users perceive them and may interact with them. Parkes et al. discuss kinetic art as an inspiration for motion design in HCI [140]. Popper provides a historical overview of the “Origins and Development of Kinetic Art” [151], where he also discusses the types of movement that are present in kinetic sculptures: the movement of the subject, the artists’ gesture, the physical movement of the sculpture, and the viewer moving in space around it. Popper stresses this spatial relationship between the viewer and the sculpture [151, p. 204]: “The space occupied by the work and the space occupied by the spectator are no longer separate. Thus a subtle game between illusory space and real space, between illusory movement and real movement, is established.” Applying this discussion to shape displays means that interaction starts long before the user touches the interface surface. It also contains notions that shape displays can express embodied gestures, and that they react to, and even guide, user movement in the space around them. Jack Burnham takes this point to a more radical conclusion, which touches on how an environment composed of computational material may one

day redefine our notions of space and time: “This is a shift from being to becoming. Kinetic works reflect this shift since kinetic works refute static space. They destroy lineal time. Kinetic works do not occupy space, they create space. Kinetic works do not contain time, they create time. Kinetic works do not interpret reality, they are reality.” (quoted in [178])

Kinetic sculptor and theorist George Rickey provides a morphology of movement, but also discusses how humans connect to physical motion on an instinctive level [162]: “Motion is measured by time, of which we all have some rather precise perception. We can compute it, sometimes with uncanny precision, witness catching a ball, passing a car on the highway, or riding a surfboard. We can measure slow-fast, long-short, pause, interval, beats per second, period of swing, coming toward us, going away from us, acceleration, vibrations separated, vibrations as a tone-these are all measurable without comparison with other objects or recollections of past experience or relation to other events in time; they have a kind of immediate measure, which, in spite of the abstractness, can give a sense of scale.” He continues with a theory of how this movement makes abstract art accessible to humans [162]: “This measure and sense of scale become, then, in kinetic art, of equivalent relevance to the human image and human scale in figurative art; even the uninitiated can bring their measure and feel that the work has something recognizable for them. This provides an entree into any non-objective art which is endowed with movement. Movement is not, in itself, esoteric; art which moves becomes accessible.” Rickey’s kinetic vocabulary has been later adapted by Moloney to develop a framework for the design of kinetic facades [124], and can similarly add to the discussion on movement in HCI. But even more interesting to me is the idea that physical movement makes that which is abstract accessible through our bodily experience. This notion can be extended from art and for instance applied to making abstract data more accessible to computer users. Schöffers’ quote [19]: “Programmed environments have the power of immediate communication” can be interpreted in a similar way. An example of how this concept applies to HCI is an early example of a Tangible User Interface, Natalie Jeremijenko’s *Live Wire* sculpture that uses the movement of a string to convey network traffic [219]




Figure 2-3: VR and Shape Displays placed on the Reality-Virtuality Continuum (graphic adapted from Milgram and Kishino [122] VR photo [130])

- a sculpture that coincidentally also bears a striking resemblance to one of the earliest kinetic sculptures, Naum Gabo's *Standing Wave*.

## 2.3 Virtual Reality (VR)

“The ultimate display would, of course, be a room within which the computer can control the existence of matter. A chair displayed in such a room would be good enough to sit in. Handcuffs displayed in such a room would be confining, and a bullet displayed in such a room would be fatal. With appropriate programming such a display could literally be the Wonderland into which Alice walked.”

Ivan Sutherland (1965) [191]

The idea of a human computer interface with the ability to control physical matter, a room in which any shape could be generated at will, was first proposed by Sutherland in 1965 as the *Ultimate Display* [191]. Building a literal ultimate display was unfeasible 50 years ago and is still far from reality today. However, rendering rich graphical, auditory and haptic sensations to trick a user into perceiving a computer-

simulated reality has become a rich research area, *Virtual Reality (VR)*. Sutherland's own subsequent work, with the invention of the *Sword of Damocles* head-mounted display [192] is a pioneering example of this approach. Recent advances in small, high-resolution displays and rendering have introduced a new generation of HMDs to the consumer market with the introduction of devices like the *Oculus Rift*, *Samsung Galaxy VR*, and *Google Cardboard*. An alternative approach to immersive VR are projection-based systems. An example is the *CAVE (Cave Automatic Virtual Environment)*, a cube-shaped room where seamless wall projections render graphics that fill the entire field of view of a user standing inside of it [104].

VR interfaces and shape displays share the same goal of reaching Sutherland's "Ultimate Display", an interface that presents information to computer users so realistically, that they can make full use of cognitive capabilities when interacting with it. They do however differ in their approach on how to reach this goal: while immersive VR interfaces attempt to substitute the users sensory experience of the real world with a simulated virtual environment, shape displays represent information in the real world, through physical objects. When placing these two approaches in the simplified representation of Milgran's and Kishino's *Reality-Virtuality Continuum* they occupy opposing ends, as shown in Figure 2-3 [122]. However, if they were to reach their eventual goal, the result would be indistinguishable to the user, and quite likely, the same interactions would apply to both of them.

## 2.4 Volumetric Displays

Volumetric displays are graphics displays where each pixel, also called *voxel*, occupies a position inside a volume in physical space [31]. This approach to rendering 3D graphics has advantages over common lenticular or parallax barrier 3D displays: As objects are rendered in real space, with full parallax and consistent vergence and accommodation cues, multiple users can freely move around the display to see objects in 3D from various angles. A commercially available volumetric display was the *Perspecta Spatial 3D Display* with a rotating 10 inch projection screen that sweeps a



Figure 2-4: Multiple users interacting with a volumetric displays [55]

volume inside a glass dome. Research by Grossman et al. investigates how users can interact with content inside the display, through touch on the surface of the display enclosure and mid-air gestures [57], and how volumetric displays support co-located collaboration between multiple users through custom rendering and interaction techniques [55] (see Figure 2-4). While volumetric displays differ from shape displays as they do not allow for the direct touch of the rendered object, they share a similar emphasis on the importance of presenting information in physical space, and how multiple people relate to a spatially rendered object as a shared reference during discussions. This research also provides us with a set of interactions for basic 3D model transformations for CAD, such as pinching gestures to translate objects, which have inspired some of our mid-air gestures.

## 2.5 Haptic Interfaces

Haptic interfaces can be grouped into two broad categories: *tactile displays*, which render tactile sensations to the cutaneous receptors in the skin and *kinesthetic displays*, which provide force feedback to the muscles by guiding or inhibiting limb movement. Another classification differentiates between *grounded* haptic devices that are attached to an external structure and *ungrounded* haptic devices, which are held or worn by the user. Haptics is a large research field that investigates multiple domains:



Figure 2-5: Haptic Workstation at EPFL [171]

new technical approaches, and applications for robotic control [109], assistive technologies [182], medical training [5] and social applications [58]. This section only provides a brief explanation of tactile vs. kinesthetic haptic devices to frame how shape displays originate from this research area and are positioned within it.

*Tactile displays* are currently the most common type of haptic device: gaming controllers and modern smartphones employ vibrating motors to enhance realism, for silent notifications and to simulate mechanical button clicks on touch screens [72]. Another tactile sensation added to touch screens is variable friction through electrovibration; an example is *TeslaTouch* by Bau et al. [6]. A user-worn form factor with vibro-tactile actuators is the *CyberTouch* glove, which renders touch sensations for VR [77]. Contact free haptic devices produce mid-air tactile sensations through a directed air vortex [185] or through ultrasonic acoustic radiation pressure [84].

While many tactile devices only produce a single sensation at the point of contact, higher resolution tactile displays exist for visually impaired users [8], pioneered by the research in sensory substitution by Bach-y-Rita [92]. Various form factors include handheld devices [13], chairs [93], and tongue interfaces [110]. Such tactile displays are composed of mechanical stimulators or electrodes that are arranged in a grid form factor very similar to shape displays, though they render haptic sensations at a

different scale.

*Kinesthetic devices* employ force feedback through a proxy device to simulate collisions with a virtual object. Examples of this approach include the PHANToM haptic interface [118] and the project GROPE Haptic Display for Scientific Visualization [16]. Kinesthetic force feedback devices may actively apply forces to the users limbs, or inhibit body movement as shown by Mitsuda et al.[123]. More recently, kinesthetic and tactile displays have been combined into devices that render both the shape and the spatial texture of objects; an example is the modular haptics setup by Fritsch et al. [40].

A common approach of haptic devices is to render the haptic sensation independent of other sensory channels, and merge them later on. Figure 2-5 shows an example of a haptic interface combined with a head-mounted display by *Immersion Corporation*. Haptics have also been combined with co-located stereoscopic graphics in projects like *The Haptic Workbench* [189]. Plesniak et al. describe the *Computational Plastic* concept, which envisions the future of real-time programmable material properties through haptics combined with real-time holography [149].

Kinesthetic haptic devices commonly render touch sensations only through limited points of contact, and therefore do not fully support active touch. Furthermore, while they provide haptic guidance and feedback, they lack perceivable affordances in their static state - the system only provides feedback while the user interacts with the device. This has motivated early research in alternative haptic devices, such as the conceptual *Robotic Shape Displays* for computer-aided design tasks by McNeely et al., where real objects and surfaces are placed at touch points through robotic actuators [120]. Such robotic shape displays would consist of *Roboxels*, cellular robots, which reconfigure themselves spatially and lock together to simulate a desired object. A similar idea inspired the *Shape Approximation Device* by Tachi et al., where a physical prop on a motorized arm moves to a contact point in space to create a touch surface for the users hand [193]. This approach of rendering real objects in space, rather than a haptic sensation tethered to the user, led to the subsequent development of shape displays [86]. In the context of haptic devices, shape displays are considered

grounded kinesthetic haptic devices. While the user perceives tactile feedback when touching the interface surface, the spatial resolution of the actuators is usually too low to actively render spatial tactile patterns. An exception to this are temporal tactile sensations - the shape display actuators can for instance vibrate when the user touches them to render additional information.

## 2.6 Tangible User Interfaces

My research in shape displays is motivated by Tangible User Interfaces (TUI) [82], which represent information and computation through physical objects that users can touch and manipulate with their hands. This guiding principle to overcome the separation of physical input and output has inspired a large body of work with various applications, form factors, and technologies [176]. Unlike graphical user interfaces, many TUI are single purpose, well suited for a specific task, but not designed to support a wide range of applications [176].

Examples of TUIs for computer-aided design (CAD) are *Illuminating Clay* by Piper et al. [147], and *URP* by Underkoffler et al., [211]. Instead of using a Graphical User Interface (GUI) to manipulate an on-screen representation of a model with a cursor, the user modifies a physical model on a desk. This approach takes advantage of the users' skills in manipulating real world objects with their body. However, utilizing physical material for modeling has a number of drawbacks, many of which are a result of the fact that physical objects are challenging to control computationally. Therefore, like many tangible CAD systems, *Illuminating Clay* follows a strict role-division: first sensing user deformations of the object for input, and then projecting graphical output back onto the object. This division works well as long as all changes to the model are formed by the users hands. However, it does not support many functions of CAD software: Operations like loading a model from a file, combining it with another model, or scaling it, would either require the user to laboriously change the physical model or to 3D print a new model for every change. Any modification by remote collaborators or a computer simulation would result in the same problem

of physical discrepancy.

To overcome these limitations, actuation has become an emerging research area of TUI [176]. A common form factor of actuated TUI is *Active Tangible Tabletop Interfaces*, which move multiple tangible objects on a tabletop display to keep them synchronized with a digital program state or a remote object. Examples of such interfaces include *Actuated Workbench* [139], *PICO* [141] and *Tangible Bots* [144].

Since their introduction, various frameworks have been introduced to analyze the properties of TUIs, frame their contributions and identify open research questions. Ullmer and Ishii propose the *MCRpd* (Model-Control-Representation physical/digital) to contrast the physical interactions of TUI compared to the Model-View-Controller paradigm of GUI [207]. Further investigating the physical mappings between objects in TUI, Ullmer et al. propose the *Token+Constraint* framework [209], which is extended into the *TAC* (Token And Constraint) Paradigm by Shaer et al. [177]. Jacob et al. introduce *Reality-Based Interaction (RBI)* to unify various post-WIMP research themes in HCI that are based on naïve physics, body awareness and skills, environment awareness and skills and social awareness and skills [88]. They discuss the tradeoffs between the cognitive advantages of RBI vs. the expressivity, versatility and flexibility of interactions that do not mimic the real-world, and advise interaction designers to find the right balance between these poles. We believe that research on shape displays is motivated by the same point: to balance between real-world interactions with physical objects and the advantages of more abstract GUI interactions.


## 2.7 Shape-Changing User Interfaces

Apart from moving tangibles on tabletops, designers and researchers have created computer interfaces that change shape to communicate information to the user. The spectrum of such shape-changing interfaces covers a large body of work with diverse form factors, ranging from special purpose devices like a curling faucet [201], to more general purpose interfaces like an inflatable computer mouse [96] or a tablet with

tilting displays [1].

The design space of actuated tangible user interfaces was reviewed by Poupyrev et al. in 2007, who also coined the term shape displays to describe the concept behind their *Lumen* project [153]. Parkes et al. have analyzed kinetic motion for organic user interfaces and propose a design language for motion in HCI [140]. Coelho and Zigelbaum have analyzed advances in material science and soft mechanics and propose how they can enable new shape changing surfaces [23]. A 2012 review by Rasmussen et al. provides a broad overview of the existing body of shape-changing interfaces and classifies related work according to the types of shape change, types of transformation, interaction, and the purpose of shape change [158].

The idea of controlling physical matter to form human computer interfaces, first proposed by Sutherland [191], has experienced a recent revival through advances in robotics and material science, such as the *Programmable Matter* concept by Toffoli and Margolus [200]. Similar to Sutherland's idea of the *Ultimate Display*, such a future interface could form any type of shape dynamically, and if it would have sufficiently high resolution, users would perceive it as a new type of shape-changing material. Researchers have proposed terms like *Digital Clay* [168], *Claytronics* [52] and *Radical Atoms* [80] for such a hypothetical computer interface. In robotics, *Bucket of Stuff* describes a similar concept of modular self-reconfigurable robots that form arbitrary shapes [226]. An example of the current state of the art in this area is the *M-Blocks* project [167], shown in Figure 2-6a. While such modular robots are impressive technical achievements, they have to overcome numerous challenges in actuation, power distribution, and networking before they will be able to create arbitrary, robust, high-resolution shapes. Therefore, current generation shape-changing user interfaces are built on simpler technology, like malleable objects with embedded actuators. These actuators can be directly embedded into the object surface in a *crust* approach [22] (see Figure 2-6b), or mounted underneath the surface to deform it externally [87] (see Figure 2-6c).



(a) Modular: *M-Blocks* [167] (b) Embedded: *Surfex* [22] (c) External: *FEELEX* [86]

Figure 2-6: Examples for different technical approaches to shape generation: modular (2-6a), embedded (2-6b) and external (2-6c).



# Chapter 3

## A Primer on Shape Displays

This chapter is intended as a starting point for researchers who plan to build shape displays. A review of previously proposed application domains, form factors, and interaction techniques leads to an analysis of the open research questions that emerge from prior work, which have motivated the research presented in this thesis.

Shape displays are defined by Poupyrev et al. as computer displays where in addition to color, each pixel can also render physical height [153]. While most shape displays do not fit this narrow definition, it has become a common term to describe computer interfaces that render physical shapes through a grid of linear actuators. Building such shape displays is a challenging process, as no commercial hardware and software platforms for these devices are currently available. Therefore, their creators often need to develop many aspects of the project from scratch, which include the engineering of the actuation mechanism, control circuits, physical design, software for rendering and sensing, and the final applications and interaction design. Such a multi-disciplinary undertaking requires a diverse set of skills, and can be daunting for researchers with limited time and resources. The design choices during the development process are driven by many factors, some of which are conflicting: one example is when the desired resolution for shape rendering is not achievable due to the size of the actuators.

What, then, are the various factors that need to be considered, how have past projects balanced these various factors, and what have been the outcomes? To better

understand this design space, this chapter reviews shape displays from a few different perspectives. A section on the chronological history points out their development over time. A hardware perspective investigates the various actuation technologies that have been explored in the past. Finally, the parameter-space discusses the factors that define how content is rendered on these interfaces. This analysis leads to our motivation for defining a set of interaction techniques in the following chapter.

Actuator arrangements similar to 2.5D shape displays exist on different scales, from fingertip sized tactile displays [8] to wall-sized displays [53]. The display orientation can vary between horizontal [114], vertical [71], or suspended from a ceiling [172]. Apart from these tethered form factors with a fixed orientation, shape displays may also be mounted on a gimbal [70], or hand-held [1], allowing users to move and reorient the shape rendering. For this thesis research and analysis of related work, we focus on horizontal tabletop shape displays with a diameter of 20 cm to 100 cm, allowing users to touch and manipulate shapes with both hands.

### 3.1 Chronological Development

The idea of rendering real objects, rather than only haptic sensations had been proposed in the context of haptics since 1993 [120, 193]. This concept inspired haptic devices based on a bed of pins, with a form factor similar to the pin screens toys invented by Ward Fleming [34], but with computer-controlled motors driving each pin. The first demonstration of such a 2.5D shape display as a computer interface was presented by Hirota and Hirose in 1995 with the *surface display*, a 4 x 4 linear actuator array mounted on a gimbal [70]. Subsequently, Iwata et al. developed *FEELEX* in 1997 with the motivation to overcome the shortcomings of haptic interfaces tethered to the users body, and demonstrated *FEELEX II* with a higher resolution for medical simulation in 2001 [86].

From 2000 onwards, researchers have been investigating various actuation techniques; while prior projects were based on electro-motors to move the shape display pins, subsequent research investigated the use of novel actuation methods to achieve

a denser spacing, or faster actuation speeds. *HypoSurface*, developed in 2001 by Goulthorpe et al., is a kinetic wall based on pneumatic actuation [53]. In 2003, Nakatani et al. presented a 4x4 pin *3D shape display* based on shape memory alloy actuators [129], and a faster, higher resolution 16x16 pin prototype, called *Pop Up!* in 2004 [127]. As part of the NSF funded *Digital Clay* research initiative, a prototype with a 5 x 5 array with custom hydraulic actuators was developed by Zhou et al. in 2005 [168]. In 2004, Poupyrev et al. designed *Lumen*, a shape display with 13x13 pins driven by shape memory alloy actuators [152] with integrated LEDs for display. In the context of this work, they also developed a variety of applications and interactions for shape displays [153]. Also in 2004, Xenotran presented the *XenoVision Mark III Dynamic Sand Table* by Xenotran, which utilized a simplified actuation mechanism in combination with pneumatic clutches instead of a single actuator for each of its 7000 pins [138, 39]. The *Gemotion Screen*, presented in 2008 by Niiyama and Kawaguchi is a pneumatically actuated flexible fabric screen with front projected graphics to display organic art [132]. *Photonastic Surface* by Oguchi et al. proposed a new method for controlling the individual actuators of the shape displays through projected light in 2008 [134].

Since 2010, the research focus has shifted to explore novel applications and form factors. Our own research, started in 2009 with *Relief* and continued with *inFORM* and *Transform*, utilizes shape displays with relatively simple actuation technology and instead focuses on novel applications and interaction techniques [114, 37]. The *EMERGE* shape display by Taher et al. uses a similar actuation approach to explore interaction with dynamic bar charts [194]. *ForceForm* by Tsimeris et al. [203] is a soft, dynamically deformable surface for medical applications [205] and tangible controls [204]. *Tilt Displays* by Alexander and Subramanian proposes a handheld form factor [1]. To support fast prototyping of applications and different form factors, Hardy et al. introduce *ShapeClip* [60].

## 3.2 Purpose and Application Domains

Most shape display projects start with an intended use case and purpose, from which the form factor, resolution, and many other parameters of the final system are derived from. For example, a GIS planning table for a group of people will be very different in size and orientation from a mobile navigation interface. So what applications have been proposed for shape displays in the past? A review of the related work reveals a vast variety of applications in the context of art, data physicalization, computer aided design, remote communication, accessibility, and functional handling of materials. As these systems are one-off museum installations or lab research prototypes, it is often not clear to what extent some of the proposed applications were implemented and tested with end users. However, by analyzing the use cases that their inventors have envisioned, we can identify the types of functionality and interaction patterns that span across different application domains.

### 3.2.1 Art

Shape displays originate from the rich lineage of kinetic sculptures (see Section 2.2), where artworks like Gianni Colombo's *Superficie Pulsante* (1959) or Ward Fleming's *pin-table display* (1979) have explored form factors to which later shape displays are quite similar. It is therefore not surprising that many shape display projects are positioned at the intersection of HCI and art. To help understand what themes these artworks follow, we can categorize them by how they interact with the audience:

- *Choreographed animations* play back pre-scripted motion, which the user views while moving around the sculpture.
- *Reactive procedural motion* is pin movement connected to an external stimuli, like a sound or gestures.
- *Interactive creatures* exhibit behavior in how they react to the user.

The boundaries between these different levels of interactivity are not strict - a sculpture may for instance switch between pre-scripted motion and interactive behavior. But the categorization helps in understanding how artists attempt to initiate

a physical conversation between the artwork and the audience through shape change. An important factor to consider is the level of intimacy for this conversation: in a setting with few people and time for engagement, more interactivity and thus a higher level of intimacy may be appropriate than in a public space.

### **Choreographed animations**


Choreographed animations do not react to user input, but play out choreographed motion in physical space. Viewers engage with the sculpture as they move around it to experience it from different points of view throughout its transformation. An example is the *Kinetic Sculpture* by ART+COM, where spheres suspended from the ceiling form patterns that morph into outlines of cars [172]. The shape transformation itself can guide the viewer to move around the sculpture by shifting between various ideal viewpoints; an example is the *Breaking Waves* sculpture by Lieberman and Washabaugh [190].

### **Reactive procedural motion**

The next level of interactivity connects the movement of the shape display sculpture to a sensor input. An example is the wave motion on the *HypoSurface* wall triggered by sounds [53], or Daniel Rozin's *Wooden Mirrors* sculpture reacting to visual stimuli like gestures [170], which invites audience participation. Movement can also be synchronized with a remote data source, like *Tele-Present Water* by David Bowen, where a grid suspended from a ceiling replicates the wave of a buoy in a remote location [14]. In this case, the projects attempt to connect the viewer to a remote location by transmitting sound or energy data in a sculptural form from the remote location to the viewer.

### **Interactive creatures**

These projects render creatures, which exhibit animal-like behavior and agency as they react to viewer input. The *Anomalocaris* application, shown in Figure 3-1a, was exhibited at Ars Electronica on the *FEELEX I* shape display. It is a creature from



(a) Anomalocaris on *FEELEX I* [86]


(b) Molebot [108]

Figure 3-1: Interactive creatures on shape displays.


the Cambrian Era, which reacts angrily when touched by exhibition visitors [86]. The *Gemotion Bumpy Screen* by Yoichiro Kawaguchi outputs abstract organic animations through a deforming membrane, onto which graphics are projected [132]. Depending on how the user pushes into the screen, the animation reacts through calm motion or pushes back with excitement. Another reactive creature is proposed for *LUMEN*, where a rendered fish responds to touching the shape display’s surface and swims towards the user’s hand [153]. *Molebot* by Lee et al. evokes the impression of a mole moving underneath the interface surface, which plays games together with the user [108] (see Figure reffig:molebot).

### 3.2.2 Data Physicalization

As defined by Jansen et al.: “A data physicalization (or simply physicalization) is a physical artifact whose geometry or material properties encode data.” [89]. The motivation behind this concept of representing data through physical objects is to allow humans to better utilize their sensorimotor system for making sense of it, which leads to cognitive benefits compared to graphical visualizations. The fabrication of these physicalizations have progressed from manual to digital tools like 3D printers, but it is currently still a slow and rather inflexible process. Shape displays on the other hand can render data as a physical shape in seconds instead of hours. They also have the advantage that the shape can change over time to represent time-based



(a) A bar chart on *EMERGE* [194].



(b) Rendering an equation on *inFORM*

Figure 3-2: Examples of Data Physicalization



(a) Northrop Grumman TerrainTable.



(b) *Tacto-phone* by Oren Horev [74]

Figure 3-3: Examples of GIS Applications

information, as well as react to user input. Past shape displays have been utilized in the context of *Geographic Information Systems (GIS)*, for the analysis of *Medical Data*, or to represent abstract information through *Physical Bar Graphs*.

### Geographic Information Systems (GIS)

An example of a *GIS interface* is the *Northrop Grumman TerrainTable*, to support decision makers in understanding terrain during mission planning and operation (see Figure 3-3a) [68]. A similar application scenario is envisioned by *Xenotran* founder Derrick Page for their *XenoVision Mark III* shape display [39]:

“ There are some obvious applications in the military area for planning and

training. Other government agencies could use the XenoVision for disaster management such as forest fires and floods. In the civilian world the XenoVision could be used for real estate development and resort planning such as golf courses and ski runs. In the entertainment market it could be used in many areas such as the display golf greens and fairways from around the world right there in the clubhouse.”

*Tilt displays* by Alexander et al. [1] has a map mode that presents terrain features through tilting displays, with the proposed use case of aiding users in planning their navigation routes. A more abstract form of representing physical waypoints for navigation is explored in the *Tacto-phone* video scenario by Oren Horev, where the back of a cellphone creates animates shapes, called *Tactons*, which the user touches while navigating through a city [74] (see Figure 3-3b).

## **Bar Charts**

The form factor of shape displays is a good fit for rendering physical bar charts, where each physical pin acts as a bar in the chart. Figure 3-2a depicts an example of such a dynamic bar charts on the *EMERGE* system by Taher et al. [194].

An important factor for bar charts is the range of the pins, so that they can render a sufficient difference between the individual values. For this reason, custom actuators with a large range have been built for systems like Regand et al., who use of physical bar graphs to render traffic data for community engagement [160].

## **Medical Data Analysis**

The *FEELEX 2* system is an interface for training doctors for palpation, where a doctor can press into a rendered organ to discover a medical condition through changes in stiffness [86]. Besides training, remote diagnosis is mentioned as a future application area by Iwaata et al. [86]. Remote diagnosis could relate to two different scenarios: when a doctor is located in a different place, and when a doctor operates inside a persons body without direct access, such as during laparoscopic surgery. An

example of adding a sense of touch to the later scenario has been proposed for tactile arrays, which have been added to the handles of conventional palpation instruments by Ottermo et al. [137].

Besides such interfaces for a physical diagnosis of simulated tissue through touch, shape displays could be used as interfaces to explore volumetric data sets, such as those produced by *Magnetic Resonance Imaging (MRI)*.

### 3.2.3 Remote Communication

Actuation is often utilized to connect the physical state of user interfaces over a distance - it can add a haptic channel to remote communication, or keep the objects of a shared distributed workspace synchronized [15]. Shape displays are an ideal platform for these concepts, as they can sense, transmit and replicate shape deformations. However, not many applications have been proposed in this domain so far, with the most elaborate example being *LUMEN*, which allows users to draw traces on top of the shape display, which are rendered as physical shapes on a connected remote display [153]. *PegBlocks*, is an educational tool for teaching physics principles, but also provides a haptic link between connected devices [146], where pushing a peg on one device will move a corresponding peg on a connected device. A proposed future application for *Pop Up!* is to combine a rangefinder with a shape display for a physical video conferencing application [129].

Beyond these initial explorations, shape displays have not yet been utilized for computer supported collaborative work (CSCW) - for instance, no example of a shared *distributed physical workspace* [15] has been created with connected shape displays.

### 3.2.4 CAD Tool

As discussed in Section 1.3, the motivation for our own shape display research stems from the CAD domain, which eventually resulted in the *Recompose* and *inFORM* projects. Other shape displays that explicitly state CAD as an application domain are *PINS* by Kaanta et al. [91], and the *Digital Clay Project* initiative [168], where

Gargus et al. proposed interaction techniques to modify CAD data through shape deformation [44]. As these interactions would have required a larger shape displays, they were prototyped with a *PHANTOM* Haptic Interface. Iwaata et al. propose CAD as a possible application domain for FEELEX [86]. *AR-Jig* by Anabuki et al. [3] is an example of a handheld CAD interface combined with augmented reality, which has inspired our own research.

### 3.2.5 Accessibility

Rendering physical shapes could give blind users the ability to experience and modify 3D models. This specific use has not been explored yet, but Poupyrev et al. have proposed to use the *LUMEN* system to overlay Braille on top of graphics as an alternative information channel for blind users [153].

This scenario is similar to how *tactile displays* render information through an array of addressable pins [215]. Tactile displays for blind users have a similar form factor to shape displays, as they also consist of an array of linear actuators. In contrast to shape displays, the pins of tactile displays usually only extend a maximum of a few mm from the surface and render binary features of braille letters and outlines, rather than 3D objects. Nonetheless, research in this area has produced concepts that could be adapted to shape displays, for instance connecting a camera image directly to tactile output in the *Optacon* system by Bliss et al. [13], or automatically converting graphical user interfaces into haptic features, proposed for the *BrailleDis 9000*, a high resolution tactile display with a 60x120 pin array and multitouch input [217, 154].

### 3.2.6 Dynamic Physical Controls

Shape displays can render physical user interface elements like buttons, touch strips and panels, to create physical affordances on demand and provide a barrier-free solution for touchscreens. This use is discussed by Iwaata et al. [86], Poupyrev et al. [153], Harrison et al. [62], and Tsimeris et al. [204]. A simplified tactile consumer device based on a similar idea is presented with the *Tactus* touchscreen overlay [21].

Besides static elements that appear and disappear on demand, tangible controls can morph between different shapes to change their physical affordances. While not implemented on a shape displays, the *Haptic Chameleon* concept by Michelitsch et al. proposes such a functionality [121]. In our own work with the *inFORM* shape display, we investigated such dynamic affordances, discussed in-depth in Sean Follmer’s dissertation [38].

### 3.2.7 Material Handling

All the previous applications in this section use shape displays to communicate with the user. However, shape change can also have a functional purpose of moving or deforming materials through shape change. This approach of using a dynamic shape mechanisms as a mold for casting materials, or as a die for forming sheet metal, has been explored in manufacturing, a process commonly referred to as *Reconfigurable Pin-Type Tooling*. Munroe et al. provide an overview of the history and technological approaches of this process [126].

As an alternative to deforming objects, *Festo* proposes a shape changing surface as a future assembly line to move objects in factory automation, an approach termed *WaveHandling*[32]. Our own research in this area has culminated in using shape displays to assemble objects from multiple blocks and has been published in the thesis of Philipp Schoessler [175] and in a forthcoming UIST paper. Shape change for moving materials is not just limited to solid objects; the *Polka Dot* system moves liquids on a water-repelling surface of a shape display for an artistic application [102].

### 3.2.8 Other Applications

*Music Performance:* The *PocoPoco* interface by Kanai et al. moves the knobs of a music controller up and down. *Education:* while educational applications could benefit from shape output, we could not find related work that proposed such scenarios.

Recently, toolkits to simplify shape display development have been proposed. These have no specific stated application, but instead allow researchers to proto-

type various form factors and applications: a hardware toolkit is *Shapeclip*, by Hardy et al. [60], which allows designers to arrange individual linear actuators on an lcd display and control them through graphics. Weichel et al. propose *Shape Display Shader Language*, a rendering framework to simplify the programming of shape displays [218].

### 3.2.9 Analysis of Application Domains

Between the different shape displays and related form factors of this review, a wide variety of applications and interactions exist. However, certain patterns are common across these different projects. Many of the shape displays use them for their advantages similar to *Tangible User Interfaces (TUI)*, i.e. that users can employ their hands to get tactile cues from the object they touch. An often proposed interaction is therefore claylike deformation through pushing down on the pins to modify the surface geometry. But our analysis of related work also reveals some key differences between shape displays and other *Tangible User Interfaces (TUIs)*:

- *Versatile Applications*: Shape displays can switch between a variety of different use cases and application domains. While most TUIs are often single-purpose form factors built for a specific application, a user may for example switch back and forth between a music player and a telepresence application running on the same shape display [153].
- *Dynamic Content*: Shape displays can physically represent changes in dynamic content. Through shape transformations, they play back time-based animations or react to user input.
- *Dynamic Representation*: Shape displays can scale, rotate and move models they render, or switch between tangible and graphical representation [153]. The combination of shape output with graphics is a common technique and allows to switch between different graphical overlays.
- *Flexible Input*: Shape displays can utilize various input modalities that are not just limited to direct touch. While users have to physically grab a static TUI object to move or deform it, shape displays can also be controlled through ges-

tures or voice commands. However, such alternative input is not yet utilized often, and early explorations have been proposed, like a flashlight tool to affect pin movement for *Photonastic Surface* [134]. Some projects also discuss potential future use cases, like capturing the users shape with a depth-sensing camera for remote communication for *Pop Up!* [129].

This comparison points at the future potential of shape displays of overcoming the limitations of static TUI, and how they can merge the dynamic nature of *Graphical User Interfaces (GUI)* with rich tangible affordances. But to fulfill that potential, they require interaction techniques that span across different applications. Conventional GUIs have well-established standards like the *Windows, Icons, Menus and Pointer (WIMP)* paradigm, which allows developers to focus on the content and application, rather than having to invent new interactions each time [69]. Such standards also provide a consistent user experience over a diverse set of contexts, so users don't have to learn new interaction techniques for different applications. Shape displays currently lack such interaction standards, which can be explained by a combination of the following of reasons:

- *Input sensing:* The input sensing capabilities vary across different shape displays, as sensing technologies constantly evolve and were not readily available until recently. An example is the recent emergence of inexpensive depth sensing cameras, which have greatly simplified the optical recognition of gestures and objects. Early shape display systems had to develop interactions for more limited sensing capabilities.
- *TUI Background:* Shape displays have a background in TUIs, which themselves are too diverse for consistent interaction standards. While frameworks to classify TUIs provide general guidelines [210, 177, 88, 33], the large variety in form factors does not allow for standardized interactions comparable to WIMP. In addition, the single-purpose nature of TUI does not require standards between different applications.
- *Market adoption:* Shape displays currently only exist as research prototypes.




Figure 3-4: Parameters for Representation: a) Area b) Pin Spacing c) Pin Diameter d) Vertical Range e) Speed f) Haptic Feedback h) Sound

Therefore, a common interaction standard may only emerge once devices become commercially available for a variety of real world applications.

Given these constraints, it may seem too early to try and develop a standardized set of interactions. Indeed, we see the main contributions of such an undertaking in stretching the boundaries of what is considered possible, and to provide guidelines and implications for future work, rather than a strict corset of arbitrary rules. As such, we identify common functionality across the various related projects and application domains, which provide us with a starting point to develop a set of interaction techniques, outlined in Chapter 4.

- *Haptic Exploration*: Touching content to explore its shape.
- *Single Touch Activation*: Touching interface elements like buttons or content to activate their functionality.
- *Deformation*: Deforming the surface geometry to modify an underlying 3D model.
- *High-level transformation*: Moving, translating and rotating content.

### 3.3 Parameter Space for Shape Rendering

As described in Section 2.7, various frameworks and taxonomies classify the general space of shape changing interfaces [158, 169, 140]. While these are well-suited for situating 2.5D shape displays within a broader context of related work, we propose a narrower, more specific parameter space to analyze the differences between specific shape display projects. By considering the unique properties of how 2.5D shape displays render content, we can also develop interactions well suited to the advantages and limitations of this form factor.

The shape displays in this thesis all have a similar general form factor: an array of pins arranged on a grid and protruding from a surface. A motorized actuation mechanism inside the system moves the individual pins to specified positions to render a physical shape on top, which resembles a bas-relief sculpture. This shape is sometimes augmented through graphics on the surface. User input is recognized through sensing embedded in the display, or through additional external sensors to capture interactions and objects above the surface. Figure 3-4 depicts the parameters for shape rendering: *Area*, *Pin Spacing*, *Pin Diameter*, *Vertical Range*, *Speed*, *Haptic Feedback*, and *Sound*. Other important parameters are: the physical interface *Surface* (see Figure 3-5), *Graphics* (see Figure 3-6), and *Tangible Tools and Objects*.

#### 3.3.1 Dimensions and Form Factor

The physical dimensions of the shape displays in this thesis are informed by the idea that a user should be able to explore and deform a shape with both hands simultaneously, but still be able to reach across the interface surface. For this reason, we designed shape displays with a side length between 15 inch and 30 inch. Besides the *Area*, the *Arrangement* of the actuators and the *Form Factor* are important. Common form factors are a square, rectangular, or octagonal footprint, in which actuators are usually arranged in a square grid or with an offset between lines to form a triangular arrangement.

### 3.3.2 Resolution

The resolution is defined by the spacing between individual actuators on the shape display. To define the ideal resolution or pin spacing for shape displays is not a straightforward task, as different touch receptors in the skin and muscles operate at different resolutions. While the ideal resolution would be high enough to render objects that are indistinguishable by touch from any naturally occurring surface, skin receptors can distinguish between surfaces with amplitude differences as small as 10 nm according to a study by Skedung et al. [183] - achieving a high enough shape resolution is therefore currently not feasible through conventional means. However, a lower resolution is sufficient for rendering distinguishable shapes that users explore with their fingers. This resolution depends if the user is moving their fingers or pressing against an objects. For tactile displays, a number of studies have been conducted to determine the ideal resolution: After a preliminary study, by Nakatani et al. conclude that a minimal pin spacing of 1 mm is desirable [128], while Garcia-Hernandez et al. recommend a pin spacing of 1.8 mm on an area of  $1\text{cm}^2$ , after finding that the discrimination performance for recognizing 2D shapes through active touch at this resolution did not significantly differ compared to exploring the shape with the bare finger [43]. Both these studies were conducted with passive props, as current actuation technology does not allow for such a high resolution.

In practice, the resolution of the shape display is not just defined by the users physiology, but also by technical considerations, such as the footprint of available actuators, the overall display size, and the cost and complexity of the interface. So from a pragmatic point of view, an important question is how low of a resolution we can get away with and still render shapes that are meaningful to the user. In our own practice, we determined that the pin spacing should be half the size of the finest feature that we would like to render or detect through the pins, so the display satisfies the Nyquist-Shannon sampling theorem. Therefore, when rendering a hand, the pin diameter should be half the size of a finger. However, downsampling and scaling content can adapt it for lower resolution displays, and physically interpolating it through

an appropriate interface surface in combination with high-resolution projection can also overcome certain resolution limitations.

### **3.3.3 Actuation Speed**


The actuation speed defines how fast the shape can transform from one state to another. This depends on the speed of the individual actuators and the bus bandwidth between the application and the actuators. Another factor is if the speed is constant vs. non-linear and if it varies depending on actuators moving up vs. down. From our analysis of related work, we determined that in order to create an interactive system without lag, a refresh rate of 15 Hz is ideal. Note that this refresh rate is much lower than common haptic force displays, which constantly need to check for collisions with virtual objects and provide feedback on these.

### **3.3.4 Vertical Range and Resolution**

The vertical range and resolution are determined by the actuation technology of the shape display. Range defines how far a pin can extend from the surface; while this value ranges from a millimeter to 100 cm, most shape displays have an actuator range of about 20 - 150 mm per actuator. Resolution can be binary, as common with solenoid actuators, or allow for hundreds of steps. In practice, resolution is often coupled to range: the longer the actuator range, the more resolution is necessary.

### **3.3.5 Haptic Feedback and Actuation Force**

When a user touches a shape, the receptors in the skin provide cutaneous (tactile) information in addition to the kinesthetic sensation from the muscles, tendons and joints. By changing the parameters of the proportional-integral-derivative (PID) controller [164] for the shape display motors, the shape may feel soft, stiff or provide dynamic feedback like vibration.



(a) Bare aluminum pins.

(b) Fabric surface and projection.

Figure 3-5: Examples of different interface surfaces on *Relief*.

### 3.3.6 Sound

While pins move, the shape display often produces sounds due to the noise of the individual motors, the sound of the mechanical linkages and the friction between the material of the pins. Apart from artistic purposes, sound is often an undesirable factor and often not discussed in related work. The sound level is usually connected with the actuation technology: servo motors are louder than nitinol actuators. It is also a function of how many actuators are moved simultaneously at what speed - when moving at full speed, our *inFORM* shape display produces up to 85dB at a distance of 1 meter, which is a considerable amount of noise.

### 3.3.7 Interface Surface

The interface surface is the actual physical material that the user touches and perceives: it can consist of individual pins, or a malleable material like flexible fabric or foam on top of the pins (see Figure 3-5). The choice of interface surface fulfills multiple purposes: a malleable material acts as a spatial low pass filter that *interpolates* between pins, and may serve as a *projection screen*, while providing *passive haptic feedback* through its material qualities like texture and elasticity. In contrast to a malleable surface, individual pins allow for more precision and higher spatial fidelity (ie. abrupt edges), and provide physical handles for pulling and pushing. Worth mentioning in this context is a very different approach of interface surfaces that are not meant to be touched, by forming ferrofluids with electromagnets, such as Kodama's *Protrude, Flow* sculptures [100], Frey's *snoil* project and *Programmable*

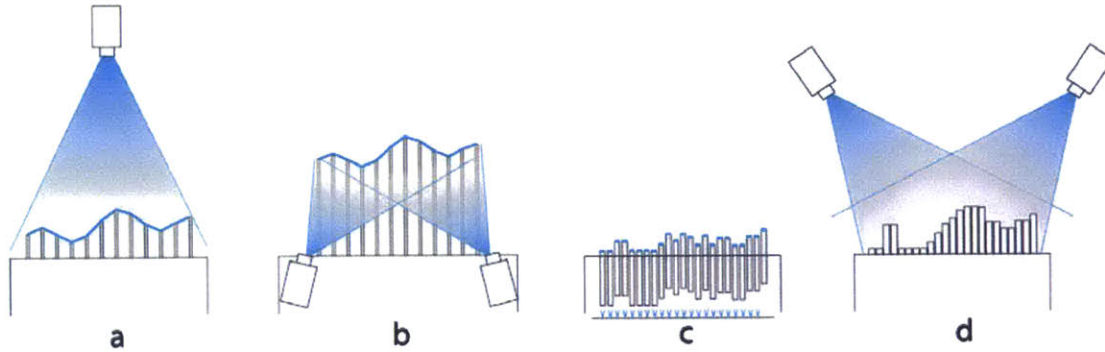


Figure 3-6: Augmenting the interface surface with graphics: a) top projection [113] b) rear projection [180] c) embedded [152] d) projection mapped onto sides

*Blobs* by Wakita et al. *Wakita2011*.

### 3.3.8 Graphics

Graphics on the interface surface add color and detail to the rendered shape to emphasize it, or to present additional information. To augment the shape with graphics, current shape displays either project information on the surface [180], or utilize embedded LED's [152]. Figure 3-6 depicts different possible arrangements of projectors and LED's for graphics. An interesting arrangement that has not yet been explored is to use individual pins as a fast moving display surface for a persistence of vision (POV) display. Such a display would function similar to current swept volume 3D displays [31], but with many individual moving surfaces, which could avoid collisions with hands and provide haptic feedback.

Figure 3-7 depicts various possible spatial arrangement of graphics displays that may complement the shape display, inspired by the display taxonomy for Augmented Reality displays by Bimber and Raskar [11, p. 72]. These include head-mounted displays, hand-held devices, optical see-through displays, graphics on shape display or its surrounding surfaces [194], or show additional contextual graphics on vertical displays. While previous work has focused on augmenting the surfaces of the shape display with graphics, we also explore the combination with *spatial 3D graphics*. Through optical see-through and video-see through techniques, we can create the




Figure 3-7: Additional graphics displays: a) head-mounted b) hand-held c) see through d) projection mapped e) vertical display

illusion of spatially co-located graphics to switch between rendering modalities (see Section 4.4).

### 3.4 Representing Content

The above parameter space is helpful for classifying the technical capabilities of different shape display hardware, but it does not discuss the various parameters that content designers can control. In this section, we investigate the design space of how content can be represented through shape displays. When defining affordances, Gibson [49] lists a number of object properties: “When the constant properties of constant objects are perceived (the shape, size, color, texture, composition, motion, animation, and position relative to other objects), the observer can go on to detect their affordances.” Rasmussen et al. characterize the design space for shape-changing interfaces as changes in orientation, form, volume, texture, viscosity, spatiality, adding/subtracting, and permeability [158]. The advantage of 2.5D shapes over many other shape-changing interfaces is that interface designers can control multiple perceived parameters at the same time: a rendered shape may appear to grow in size

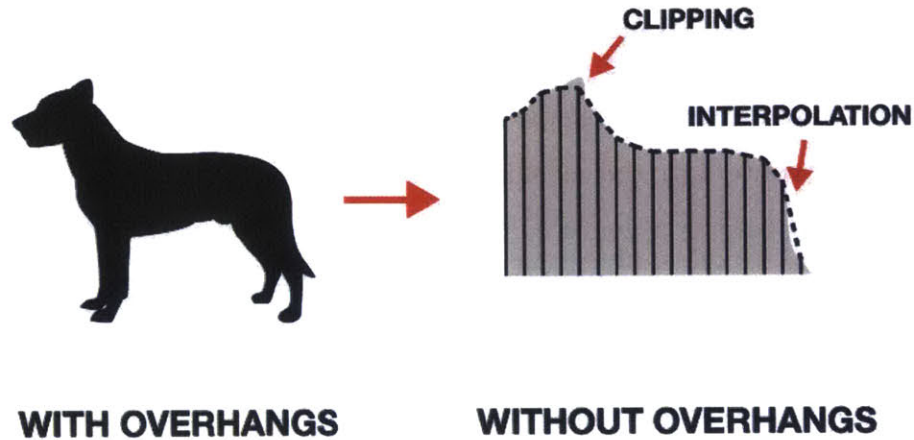


Figure 3-8: Rendering a 3D model on a shape display usually results in geometry with missing overhangs, and clipping and interpolation errors.

while moving across the table and changing its color.

## Shape

The range of possible shapes is defined by the degrees of freedom of the individual actuators. On 2.5D shape displays, each pin can move along one dimension, by extending and retracting from the table surface. When rendering bar charts, like proposed by Taher et al., each pin can be mapped to a data point [194]. Geospatial data, where each point for a given latitude and longitude has a distinct elevation, can also be easily mapped to pin height. In such cases, the designer only needs to consider the range of the values and if they can be meaningfully represented by the shape display or need to be scaled or cropped to fit. More complex 3D models are less straightforward to represent, as rendering them in 2.5D will likely result in loss of information. An example is shown in Figure 3-8, where the original 3D model is rendered without overhangs and clipping and interpolation errors are introduced due to the limited resolution.

## Position, Size and Orientation

While the individual pins of a shape display can only along one dimension, the output dimensions for content can be considered differently. The user perceives adjacent pins on the shape display as a whole connected shape - a cognitive process explained by the law of closure in Gestalt theory [101]. The perceived physical shape can move and deform along more dimensions than individual pins would allow, enabling geometry that rotates, moves and scales arbitrarily along three axes. This effect is similar to graphics on a low resolution bitmap display: as adjunct pixels light up, viewers recognize movement patterns like for instance render scrolling text. In many ways, this flexibility makes interaction design for shape displays compelling for interaction design, as it allows for a much larger design space than other shape-changing interfaces with confined, pre-determined transformations. We discuss this space for interactions in the following chapter.

## Motion

Physical motion is necessary for shape change, and therefore all shape displays have an inherent kinetic design element to them. But how that motion is carried out can vary greatly: it can be based primarily on a functional approach, where all actuators move at their maximum speed to render the target shape, or it can follow a carefully choreographed sequence with deliberate timing. Early on in our own experiments with motorized shape transitions, we observed that fast moving parts and abrupt changes in direction would startle or even frighten users and the speed of shape transitions emerged an important design element. But motion design can be more than a necessity to avoid startling users: it can carry meaning, convey information and evoke an emotional response. Vaughan compares motion in HCI to choreography in dance, and proposes to define it through path, volume, direction, and velocity [213]. Rasmussen et al. adapt this definition for their analysis of shape-changing interfaces [158]. Parkes et al. propose a design language for kinetic organic interfaces, by discussing how the interface *form and materiality* affects motion, how

functionality like *kinetic memory* can be utilized in the design process and how adding noise to parameters like *repeatability and exactness* leads to more organic motions [140]. Djajadiningrat et al. discuss the importance of movement from the perspective of industrial design and introduce the term of *4D form* for dynamic, shape changing objects [26].

### **Stiffness**

The stiffness of an element determines if a user can deform it, and varying stiffness across a rendered shape may guide the user in how to deform it. In addition, stiffness may change while a user presses into an object, to simulate structures inside of it. To produce this effect requires a high bandwidth, comparable to other haptic force displays. Another haptic tactile dimension can be added through vibrations with small amplitude and high-frequency that the user feels when touching the shape.

### **Visual Texture and Physicality**

Color and visual texture can augment the rendered shape with graphical details, or present an alternative layer of information, as proposed by Poupyrev et al. [153]. Beyond graphics on the interface surface, the rendered shape can be combined with spatial 3D graphics. These graphics can be combined with the pins retracting and extending rapidly from the interface surface, so a shape may appear or disappear, to control their perceived physicality. By combining the shape disappearance with spatial graphics, objects may seem to transition between a tangible, solid shape and a gas-like virtual state, as discussed in detail in Section 4.4.

### **Additional Tangible Objects**


In addition to rendering shapes directly through their surfaces, the shape displays can also convey information indirectly, by applying force to objects resting on top of them. A shape display can, for instance, move a sphere along an axis to represent time, or the movement of objects on top can be linked to connected, distributed objects.




# Chapter 4

## Interaction

Interaction with shape displays is strongly coupled to the strengths and limitations of the physical rendering hardware, like the fact that most of them are able to sense input by pushing down on the pins, but cannot sense when users laterally press against a rendered shape. As a language for interaction emerges when analyzing and developing different applications, designers have to decide is if those interactions should be developed around the limitations of current hardware or incorporate new sensing techniques and interaction modalities. We chose to follow the second route, as we see a rich opportunities of interactions that are not constrained by embedded sensing and by interaction techniques constrained to direct touch. Our interaction prototyping is guided by the following principles:



(a) Touching the surface.




(b) Deforming a model.

Figure 4-1: Interactions with bare hands through touch and deformation.



(a) Mid-air gesture.



(b) Tangible object.

Figure 4-2: Interactions beyond direct touch: gestures and tangible objects.

- *Input Hardware:* To overcome the limitations of embedded sensors, we utilize hardware that allows for the quick prototyping of various interactions without having to integrate or build new sensors for each one of them first. For this reason, we resort to depth cameras to allow for detection of mid-air gestures, touch and tangible objects. This choice prevents us from developing mobile form factors and has limitations in terms of precision and occlusions, but we can compensate for these in a lab setting and we anticipate that it will get replaced with more specific sensor hardware once successful interaction techniques have been identified.
- *Scope:* We define our interactions for a specific form factor: shape displays that are horizontal, allow for multiple users to stand around them, are larger than a single hand, but small enough to reach across them. While some of our interaction techniques may be adopted to mobile or wall-sized shape displays later on, we do not try to design interactions for all of them, to avoid interactions that are too generic to be meaningful. The tabletop form factor has the advantage that we can utilize gravity with objects on top of the shapes.
- *Generalizability:* To ensure that our interactions scale well across various applications, we base them on the analysis of different application domains in the previous Chapter (3.2), which provides us with a large sample of common functions. While the interactions may change as commercial systems will be adopted by the market, these functions serve as a generalizable reference.

As discussed in Section 3.2.9, an analysis of prior work identified the following common interactions that past shape displays explored:

- *Haptic Exploration*: Touching content to explore its shape.
- *Single Touch Activation*: Touching interface elements to select them or activate their functionality.
- *Deformation*: Deforming the surface geometry to modify an underlying 3D model.

In this analysis, higher-level spatial transformations were strikingly absent. Commands to move objects were limited to touching content to select it, then touching somewhere else to signal that the content should move to that target location. But spatial interactions like swiping, pushing or pulling to scale, rotate and translate content were absent - likely due to interface constraints like the uneven surface and lack of lateral input sensing. Some projects, like *EMERGE* resort to richer interactions like swiping on the flat surface surrounding the table [194]. To explore such interactions, we propose the following functionality:

- *Selection of objects or subset of geometry*: The ability to quickly select one or more objects or geometry
- *Transformation of the selected geometry*: Translating, rotating and scaling the selected object or geometry

We argue that while these interactions are common building blocks that diverse applications would benefit from, they have not been supported in much detail is because of the limitations of direct touch and deformation on shape displays.

## 4.1 Input Modalities

### 4.1.1 Touch

Touching the interface surface allows users to experience the shape of objects through their hands, as shown in Figure 4-1a. Beyond experiencing the shape, a detected




Figure 4-3: Limitations of direct touch interactions with shape display.

touch may also trigger software actions, like selecting an object, move it, paint its surface or annotate it. These touch interactions are similar to interacting with multi-touch displays, though shape display application designers have to deal with an added ambiguity: While users of touchscreen displays usually only touch the screen when they want to interact, shape display users may touch a shape to haptically explore it, while also accidentally triggering a touch event. Another limitation of direct touch is the irregular interface surface, which may prevent the user from physically reaching certain parts they mean to touch (see Figure 4-3a).

### 4.1.2 Deformation

Beyond touching the surface, the user may press into a shape to deform it. Sensing pressure input can resolve some of the ambiguities of touch interaction, and also add a haptic dimension to the interaction. Designers may for instance program the surface to simulate a malleable, claylike material for modeling operations in a CAD application (see Figure 4-1b), or to feel like rigid objects that move as a compound shape when pressing onto them. Pulling pins up to extrude shapes from the surface adds another dimension to deformation. However, the interface surface may not afford pulling as it does pushing, as shown in Figure 4-3b.

### 4.1.3 External Controllers

To resolve the ambiguities of direct touch and deformation, we experimented with additional input modalities by connecting a standard trackball mouse for global surface navigation. A series of interactions map to the features of the trackball: rolling the ball in any direction translates objects on the surface accordingly. Similarly, pressing either mouse button activates scaling commands.

In order to compare the use of an external controller against directly touching the shape, the input features of translating and scaling were also implemented through pressure input. When pressing an edge pin on the interface, shapes are translated from the center in the direction of the pin. Zooming in is achieved by pushing on center pins, while zooming out by pushing on two opposing edge pins. Upon comparing the two implementations, each one presents a tradeoff: the external controller reduces interaction ambiguity, it also it removes the locus of interaction from the table, thus requiring the user to shift attention between the shape display and the external controller, interrupting the flow of interaction.

### 4.1.4 Mid-Air Gestures

While touching and deforming the shape, the users hands are constrained by the degrees of freedom of the actuators. Learning from our previous tests, we propose to add gestures as an interaction modality for a 2.5D shape display (see Figure 4-2a). Gestures above the surface provide additional degrees of freedom for input, while maintaining a clear causal link between the users hand movements and the resulting shape actuation and without shifting the locus of interaction towards an external object. By expressing gestures directly above the surface, a user can seamlessly switch between selection, manipulation, and translation of objects on a 2.5D shape display. When coupled with direct touch, a user is able to maximize functionality without creating input ambiguity. A grammar of gestures has been implemented to explore basic functions used to interact with a 2.5D shape display. We have found that the most fundamental set of gestures includes: selection of an area, translation of the




Figure 4-4: Mid-air gestures for selecting and translating content.

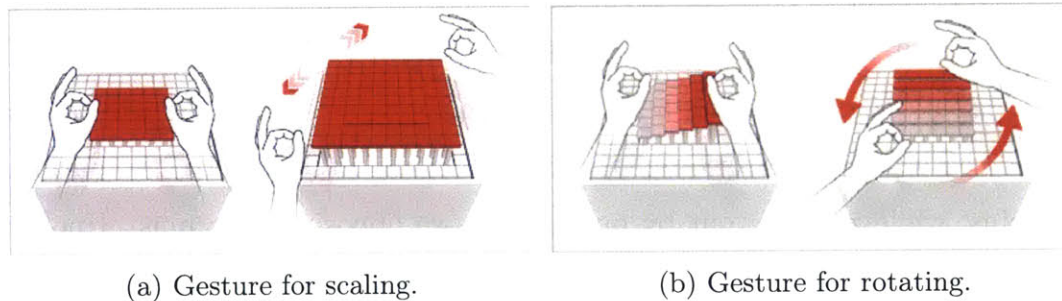


Figure 4-5: Mid-air gestures for scaling and rotating content.

selection, rotation of the selection, and scaling of the selection. Further description of these techniques follows:

*Selection:* In order to select a subset of the surface the user forms two parallel vertical planes with their hands (see Figure 4-4a) - a gesture commonly used to indicate the dimension of an object (e.g. “I caught a fish this big!”). The system indicates a selected area with a projected square selection rectangle. A two-finger pinch on either hand locks the selection dimensions, enabling manipulation through a number of gestures. These are based on the pinch gestures introduced by Wilson [222]. When pinching, the user forms a mental connection between the shape underneath and the location of the gesture.

*Translation:* By adjusting hand position along the X, Y, or Z-axis, the user can simultaneously manipulate the height and position of the selection (see Figure 4-4b). After reaching the desired height and position the user can release the pinch gesture, saving surface state, and resetting the interaction state back to selection mode.

*Rotation:* By rotating the locked hands about the X, Y or Z-axis, the selection rotates accordingly (see Figure 4-5b).

*Scaling:* By changing the distance between the locked hands, the selection scales proportionally (see Figure 4-5a).

### 4.1.5 Tangible Objects

While mid-air gestures add expressivity, they lack tangible feedback for the user. We therefore explore objects on top of the shape display to interact with content. A particular focus is on how with the physical geometry of these objects interacts with the interface shape. The rendered shape can constrain objects to guide the user, similar to the Token+constraint framework proposed by Ullmer et al. [209]. Unlike previous work, these constraints are able to dynamically change to reflect underlying software state changes [37].

Besides passively constraining the user from moving an object, the shape display can also apply mechanical force to move it in a variety of ways. An example is when the shape display forms a slope that a ball rolls down on. This technique for actuating passive objects does not require any special material like magnets or robots, which are commonly required for active Tangible Tabletop Interfaces [37].

## 4.2 Dynamic Scale and Modality

The above interactions for content transformations can be utilized to represent and navigate larger, more complex data sets. Our design goal for shape displays is to create an interface that combines the best of two worlds: the expressive qualities of physical materials with the powerful digital tools of computer-aided design (CAD) software. These tools allow us to parametrically change model geometry through control points, constraints, and global parameters. Another important aspect Graphical User Interfaces (GUI) for CAD is the fluidity of rendering the model; i.e. the ability to change how information is presented to the user at any given moment. CAD heavily relies on continuous adjustments to the viewport and the rendering style to allow the user to modify a model from various angles and at different scales. Figure 4-6 shows an example in the *3ds Max* environment, where a scene is rendered from

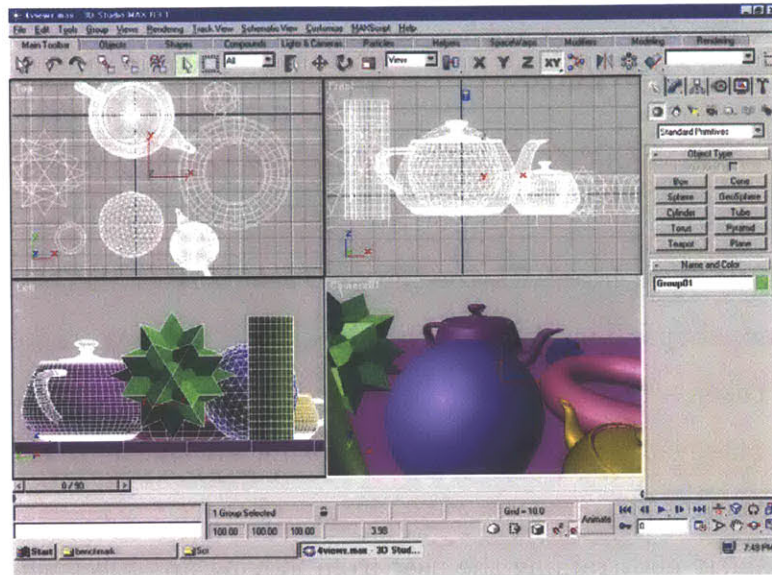


Figure 4-6: Multiple viewports and rendering styles in modern CAD software.



Figure 4-7: Building a physical architectural model. (Image courtesy by Denise Coté.)




Figure 4-8: Interaction Models of GUI and TUI (adapted from [207])

multiple viewports and in different styles. This functionality is necessary due to the shortcomings of working on a 2D display: being able to zoom, pan and rotate the view is crucial, as users cannot simply move around the model in physical space to change their view, like a craftsman would when working with a physical model, as shown in Figure 4-7. But beyond addressing display limitations, this ability to work on models at various scales from a city level down to small details, and to effortlessly switch between different model representations like solid and wireframe, far surpasses the capabilities of traditional physical model making.

The underlying GUI functionality to, on one hand, modify a digital model and, on the other, separately control how it is rendered on a screen, is reflected in the *Model-View-Control* interaction model (MVC), which was developed for Smalltalk-80 (see Figure 4-8). An input device, like a mouse, acts as a *control* to modify the underlying data *model* and to change the *view* onto that model, which is rendered on a graphics display. Ullmer and Ishii have adapted the MVC model to describe Tangible User Interfaces (TUI) through the *Model-Control-Representation (physical and digital)* interaction model (MCRpd) [207] (see Figure 4-8). The underlying data *model* of MCRpd is similar to the one in MVC, but instead of a view, it is accessed through a physical and a digital representation. The *physical representation* is a real, physical object that represents the model and also acts as a *controller* to modify it. The *digital representation* is an intangible output, such as graphics or sound, and

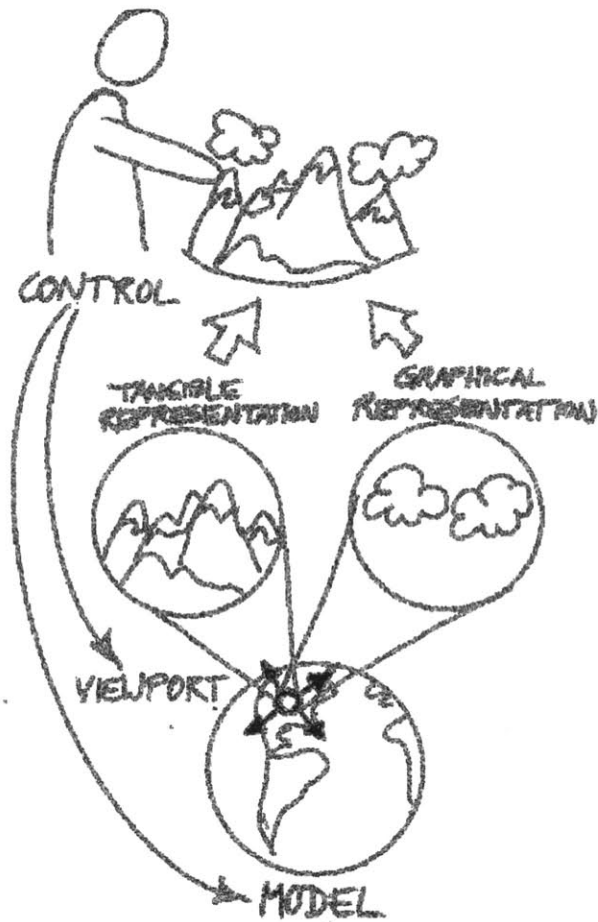


Figure 4-9: Concept sketch to show the viewport and the modality of the dynamic tangible representation.

is closely coupled with the physical representation. As the two representations are closely linked, changing the view or output modality is not as straightforward as it is in a GUI. It is also less critical, as users can move their body in space to get a closer look at an object or view it from different angles. For these two reasons, many TUIs only support a fixed scale and modality of representation, and changing them is not discussed in the MCRpd framework.

While past shape display research has not investigated the concepts of dynamic physical scale and modality, I argue that these platforms are ideal for exploring these concepts in the context of TUI - rendered models can grow, shrink, rotate, translate and disappear programmatically, while co-located spatial graphics complement these

transformations and allow to switch between physical and virtual representations. Figure 4-9 depicts this concept of dynamic scale and modality, where the interface represents a slice of the earth.

Besides manipulating content, the user can also change the scale and portion of the physical representation through *ShapePorts* and switch between or mix tangible and graphical representations through *Sublimate*. The following sections explain these concepts and the resulting interaction techniques, as well as example applications we developed to demonstrate them.

### 4.3 Dynamic Scale with ShapePorts

To support dynamic scale, we propose *ShapePorts*, a tangible adaptation of the well known viewport concept in computer graphics for rendering 3D scenes to a graphics window [35, p. 210]. In information visualization, Herr and Shneiderman refer to “manipulable viewports onto an information space” as one of the fundamental building blocks for exploring datasets [65]. The short film “Powers of Ten” by Ray and Charles Eames is a profound demonstration on how continuous zooming can convey the universe at different scales, and thus provide a better understanding of it [20]. Perlin and Fox apply the concept of continuous view transformation to interact with a large information space with the Pad prototype [145].

The concept of a *ShapePorts*, where the physical component of a TUI morphs to represent various portions of a model or data set has not been previously proposed. But there are numerous examples of TUI to manipulate a graphical viewport through spatial input: In the *Mixed Reality Interface* by Uray et al. [212], a physical camera prop on a tabletop map represents the virtual camera object of a 3D scene rendered on a separate display. Moving the camera prop on the map changes the corresponding viewport. On tabletop displays, physical handles can scale and rotate the underlying map view, such as in MetaDESK by Ullmer et al. [208]. PICO by Patten et al. builds on this interaction, and computationally moves physical objects on the tabletop display to maintain a correct relative position on the map view [141]. Spindler et al.

propose *Tangible Views for Information Visualization*, which are lightweight spatial displays to explore data sets in relation to a large tabletop display [187].

Related to the concept of a ShapePort is the *Tangible Lens* concept, where a TUI on a tabletop display will change shape as it is moved over content. Examples are the *Haptic Tabletop Puck* by Marquardt et al. [117] and *G-Raff* by Kim et al. [95]. More recently, Hardy et al. have proposed to utilize this concept to prototype shape displays with *ShapeClip* [60]. Taher et al. propose techniques to render, navigate and annotate dynamic physical bar charts [194], with a framework based on the taxonomy by Herr and Shneiderman [65]. We propose a more generalized approach to ShapePorts, which are not only specific for bar charts, but can scale across various application domains.

In software, the ShapePort is rendered similar to a graphical viewport: when rendering a scene in *OpenGL*, the color output is sent to the graphics display, while the depth buffer is rendered into a separate texture and sent to the shape display. The two main questions from an interaction design perspective are how to navigate the viewport, and how to provide context, so the user is aware of the larger model, and what part of it is currently represented.

### 4.3.1 Interaction Techniques

Figure 4-10 depicts the possible transformations to the ShapePort. As we do not utilize shear, the representation can be scaled, translated and rotated in 3 dimensions, which makes for a total of 9 degrees of freedom for navigating data sets. However, depending on the type of data model and application, it is often useful to restrict the ShapePort navigation to fewer degrees of freedom:

- *Scaling*: Scaling the viewport grows or shrinks the window: We usually only scale uniformly - except for when we want to fit content in the vertical dimension.
- *Translating*: For planar models, translations is only in 2 dimensions. For 3d objects, often only rotation is employed instead of translation.




Figure 4-10: Possible viewport transformations.

- *Rotating:* rotation only in one dimension for planar models.

### 4.3.2 Providing Context with Graphics Displays

While interacting with a subsection of a model through the tangible representation, the user needs to be aware of where this representation is situated in the larger context. There are different techniques to provide this context to the user. Interaction techniques to transform the viewport can help to gain an overview, by for instance briefly zooming out to the larger context and back in to the detailed view. This way, the user builds up a mental model of the larger dataset and where the current subset is located within it. In information visualization, various methods aid navigation by simultaneously rendering an overview along with a detailed view [65]:

- *Focus plus context:* The detailed view is situated within a surrounding contextual view, which contains a less information and helps users to orient themselves [7].
- *Overview and detail:* An overview map is rendered next to the detailed view.




Figure 4-11: Combining a ShapePort with a surrounding tabletop projection and wall displays for context.

This map depicts the larger dataset and highlights, where the current view is located within it.

- *Distortion techniques:* A common distortion technique is the fisheye view, where the detailed view is surrounded by a distorted context view, which contains gradually less information [41]. According to Heer and Shneiderman, while intriguing, this technique may however disorient viewers [65].

We appropriate these techniques to augment the detailed tangible representation with a graphical contextual view. Figure 4-11 shows an example, where the ShapePort of an alpine landscape is augmented with a projected context view on the table, and a 3D overview on the wall displays.

### 4.3.3 Applications

#### CityScope

*CityScope*, a is a GIS application for urban planners on the *inFORM* shape display, implemented by Sheng Kai Tang and Yusuke Sekikawa. It demonstrates a combination of a ShapePort with graphical contextual screens 4-12. The shape display renders





Figure 4-12: CityScape setup: shape display with a wall display rendering a contextual view.



(a) Rendering buildings.

(b) Rendering population density.

Figure 4-13: Various data sets rendered in CityScape.

a portion of a GIS data set with information layers like buildings (see Fig. 4-13a) and population density (see Fig. 4-13b). The user can select the information layer and navigate the data set through gestures. A large contextual wall display shows an overview of the urban map and highlights the portion that is currently rendered on the ShapePort – this highlight is a white wireframe box and can be seen in Figure 4-12.

## 4.4 Dynamic Modality with *Sublimate*

“As physical elements, TUI artifacts are persistent – they cannot be spontaneously called into or banished from existence.”

Ullmer and Ishii [207]

A defining element of TUIs is the persistence of physical objects, which unlike graphics do not disappear from one moment to the next. Our goal with the *Sublimate* project was to challenge this assumption. As discussed in the previous section, the ability to change the representation of an object provides a powerful tool for interacting with information, for instance by switching from a solid to a wireframe model in a CAD application. Recent research in 3D user interfaces allows us to represent physical objects with a similar flexibility, by combining actuated tangible shape displays with spatial 3D graphics. We seek to find the hybrid of the two by thinking about physical material density as a parameter in 3D rendering. We want to explore, how both digital models as well as handles and controls can be rendered either as virtual 3D graphics or dynamic physical shapes, and move fluidly and quickly between these states. Figure 4-14 depicts an illustration of this vision, with a user interacting with the model of a brain that is half physical and half virtual.

Our vision of *Sublimate* is a human computer interface with the ability to computationally control not just the shape, but also the materiality of physical objects. An object rendered through this system can rapidly change its visual appearance, physical shape, position, and material properties such as density. While such a sys-



Figure 4-14: Concept sketch of a user interacting with a partly physical model on the *Sublimate* system.




Figure 4-15: The transitions between *spatial 3D graphics* and *shape output* are inspired by phase transitions.

tem does not currently exist and might be physically impossible to build even in the future, we can build interfaces which appear perceptually similar to the user through a mix of actuated shape displays and spatially co-located 3D graphics. In this section, we focus on computationally controlling a specific parameter: the physical density of objects. Objects rendered through our system can rapidly switch between a solid physical state and a gas-like floating state. With programmable affordances, objects can be physically rendered when needed and are still visible when not. We call this concept “Sublimate”, as it is inspired by the phase transition from solid to gaseous in a thermodynamic system.

The thermodynamic phases of materials that humans interact with on a daily basis are solid, liquid, gaseous (while plasma is actually the most abundant form of ordinary matter in the universe, it is less commonly handled by humans). Material properties like density rapidly change between these phases, as one can easily observe in ice, water and steam. We apply this metaphor to the relationship between physical and virtual output in a *Sublimate* interface (see Figure 4-15). We are inspired by state transitions, particularly from a solid to a gas (sublimation), and a gas to a solid (deposition). We use this metaphor of state change as a way of thinking about interaction between shape output (solid), 3D virtual graphics (gas). Any object in

the sublimate interface can be rendered through a shape display (solid phase) or through spatially co-located 3D graphics (gaseous phase). We call the transitions between shape output to virtual graphics “sublimation”, and the transition between virtual graphics to shape output “deposition”. But we can also render objects that are partially sublimated, allowing the user to touch certain portions of an object while others remain intangible.

While past research has combined haptic interfaces with co-located 3D graphics, the role-division between these two modalities has been very distinct [173]. Popular haptic interfaces such as a *PHANTOM* render haptic sensations for separate points and do not allow users to touch and manipulate objects with bare hands. For this reason, users of current VR systems are commonly aware of the fact that the represented object is not real. In contrast, actuated shape displays form real physical objects, which the user can see, touch and manipulate with bare hands. However, current shape display projects trade the advantage of real materials for the flexibility and realism of high-resolution graphics present in VR and AR interfaces. We propose to combine these two modalities to create a rich computing experience neither one of them could provide on its own. Particularly interesting to us are interactions enabled by the ability to rapidly switch between rendering content as a solid object in one instance and intangible floating graphics in the next.

Our vision is that any information can be rendered in real space as virtual graphics or in physical form. We believe that the most interesting aspect may not be either state alone, but rather the fast transition from the virtual to the physical, and visa versa. Thus we are not only interested in augmenting shape displays with co-located virtual graphics, or adding haptic feedback to augmented reality, but also how the transition between physical and virtual can be utilized to improve interaction. Opacity has been an important tool in graphics and user interfaces, how can we add another axis of “physicality”? Parts of physical models can be replaced by floating graphics, allowing the user to physically manipulate a part inside. Virtual sliders can become physical when they need to be touched or modified. It is this rich area of state transition between physical and digital that we seek to explore.

In order to explore this space of Virtual/Physical State Transitions we designed two implementations of a system called Sublimate, which combines spatial augmented reality with actuated shape displays. The first combines a situated AR display, utilizing a half silvered mirror, stereo display and head tracking, with a shape display to co-locate 3D virtual graphics and a 2.5D physical surface. The second uses a tablet based AR system to add virtual graphics to the scene. Both systems allow for direct interaction from the user, both through mid-air interaction with a wand and through physical manipulation of the shape display. We also outline compelling use scenarios for augmenting shape displays and for state transitions.

The guiding principles for the design of a Sublimate system are:

- The output should be perceptually as close as possible to real world objects. This means that instead of solely providing a haptic sensation for limited points of interaction, the aim is to render real objects. Users should be able to touch these objects with their bare hands and naturally view them from from different sides.
- Synchronized output channels, with the ability to rapidly switch the rendering of an object between them. The system can represent information as graphics, physical objects or both.
- User input through multiple modalities. Users can interact with the system through symbolic commands, gestures and direct touch, based on the physical channel they currently interact with.

#### 4.4.1 Related Work in Augmented Reality

To support individual and decoupled control over an object’s visual appearance and physicality, we need two different techniques. First, we need a display technology that can both show graphics floating in mid-air, as well as overlaid and registered with a physical object. Second, we need techniques that allow us to control the presence of an object’s physical parts or components.

Since Ivan Sutherland’s vision of the *Ultimate Display* [191], researchers have

aimed to create an immersive environment allowing for virtual and physical elements to be rendered anywhere in 3D space. Although there has been much research in 3D user interfaces, from Virtual Reality (VR) using head mounted displays or CAVEs to Augmented Reality (AR), there has been less work on rendering physical forms. The most popular systems have used articulated arms (PHANTOM) [118] to provide haptic feedback, yet these often only allow for a single point of interaction, or require the user to be instrumented with heavy gloves or cables. These devices have been combined with spatially co-located 3D graphics through VR and AR systems to provide haptic feedback and direct input [173, 150]. We believe these systems still fall short of the vision of the *Ultimate Display* as they cannot physically render the forms, but rather only provide haptic feedback at discrete, limited points of contact. Another thread of research emerging from tangible UIs has focused on trying to render physical forms; an example is actuated shape displays, which can render physical surfaces through a series of actuators [114]. We see the combination of these trends, 3D virtual graphics and actuated shape displays open up a rich area of research.

### **Situated see-through displays for spatial AR**

Numerous projects explore techniques where a partially transparent, see-through display augments real objects or environments with superimposed information or graphics [11, 10, 135]. These Spatial AR systems can also be combined with passive tangible input in the real world [66].

### **See-through displays with co-located manual input**

Schmandt [174] describes an early setup, which emphasizes the perceptual advantages of co-locating stereoscopic imagery with the user's hand and input device. A half-silvered mirror is used to reflect 3D graphics, which is optically merged with the user's hands underneath, registered using a magnetically tracked 3D input device. although there are other techniques [227]. Toucheo [59] demonstrates how these configurations can be combined with multi-touch surfaces and on-surface interaction techniques for 3D manipulations, while HoloDesk [67] uses depth cameras to explore whole-hand

interactions, object tracking, motion parallax and physics simulations for enhanced realism.

### **3D displays with co-located tactile feedback**

Co-located setups can also be extended to incorporate tactile feedback through haptics. The Haptic Workbench [189] adds single-point force feedback through a PHANTOM device, a configuration also explored by Scharver et al. [173] in an immersive interface for tangible design of cranial implants. Plesniak et al. [150] describe the Computational Plastic concept, which envisions the future of real-time programmable material properties through haptics and real-time holography. They demonstrate a number of proof-of-concept systems based on single-point haptics and holograms. Touchable Holography [75] enable force feedback without mechanical devices by using ultrasound for a tracked finger in a 3D display.

### **Projection-based AR**

Projection-based approaches have been explored in many projects to alter the visual properties of physical objects [157], particles [148], surfaces [9], or the user's body [61]. One motivation for such systems is that they can modify appearances without additional surfaces, materials or user-worn equipment.

### **AR interfaces for control of physical objects**

AR is also well-suited for visual support and feedback during control, manipulation and actuation of devices and objects. Tani et al. [199] describe a user interface for manipulating physical controls on remote machinery through an augmented video interface. TouchMe [64] applies direct-manipulation techniques for remote robot control using video see-through AR and a touch-screen interface. Ishii et al. [83] enable real-world pointing and gesturing for robot control, using a tracked laser, where visual feedback is provided from a projector.

## Shape-changing AR interfaces

AR-Jig [3] is a 3D-tracked handheld device with a 1D arrangement of linear actuators, which is used for shape deformation and display, as the device is moved in space and interacts with virtual geometry that is viewable through an AR display.

The typical motivation for both AR systems and shape-changing interfaces is to unify virtual and physical representations to enable richer interfaces for viewing and interaction. Work like AR-Jig has explored how to co-locate haptic feedback with AR. With this paper, we introduce additional expressiveness, by enabling dynamic variation of the amount of graphics and physical matter used to represent elements in the interface, and exploring state change between rendering either as virtual or physical output.

### 4.4.2 Interactions

“There is no reason why the objects displayed by a computer have to follow the ordinary rules of physical reality with which we are familiar.”

Ivan Sutherland (1965) [191]

As the modalities of shape output and virtual graphics are synchronized, the system can render an object in either one of them independently, or in both modalities at the same time. We propose how transitioning from one modality to another can not only enrich existing interactions, but also enable entirely new ones.

#### Physical to Virtual: Sublimation

As previously described, the Sublimate system can render physical shapes. However, there are many conditions where users may desire to interact with a purely virtual version of that data or control unit:

- Reaching through physical objects: physical shapes can disappear while leaving intangible graphics behind. Users may need to physically manipulate some

object inside of a larger object, or behind an object; sublimation allows the user access to these regions quickly.

- Increasing degrees of freedom for control elements: Virtual controls are not constrained by the degrees of freedom of the shape display. The user may want to switch from precise constrained interaction with the physical control to unconstrained mid-air interaction with graphical feedback.
- Observing space and safety constraints: If objects are too large to render as physical shapes or their physical shape would be harmful to the user, they can be rendered graphically instead.
- Visualizing impending physical output: Rendering a graphical preview before the physical shape can inform users of impending consequences and allow them to cancel or confirm the output. Pre-visualized physical transformations are less startling to the user, which is particularly useful if the generated shape will touch the user or other physical objects in the space.

### **Virtual to Physical: Deposition**

The Sublimate system also allows for interactions driven by the transition from virtual to physical:

- Having physical affordances of user interface elements appear when needed: sliders, touchpads, handles and buttons are rendered graphically and appear physically when needed. Or a certain contexts or tools can trigger Deposition; for example when a user brings a pen tool to a virtual surface, it becomes physical to allow for annotation.
- Objects are able to morph to different shapes to adapt to changing interaction constraints: An example would be rendering a physical surface of a virtual object when touching it, then flattening the surface as a pen tool approaches. This would allow for the user to experience the shape through touch and conveniently draw on it with a tool.
- Restrict movements or constrain interactions where prohibited: physically rendered barriers can block access to certain areas - for example users "freezing" a

data set so no one changes it when they are gone.

### **Mixing Virtual and Physical**


In addition to transitioning between states, there are many interactions that benefit from the combination of shape output and graphics. Like in classic AR applications, spatial graphics can augment the physical shapes to provide additional information to the user. An example is visualizing the wind flow around objects. Graphics can also guide input when learning to use the system, by rendering spatially overlaid help and tutorials. Another application is to overlay alternate versions of an object onto its physical shape in CAD scenarios, similar to “onion skinning” in animation software.

In addition, graphics can help to compensate some of the limitations of current generation shape displays. They enhance the visual resolution, size and scale of shape output, and augment features that 2.5D shape displays are unable to render, such as overhangs.

#### **4.4.3 System Setup**

We built two proof-of-concept setups to prototype the envisioned interactions of the “Sublimate” concept. Each setup consists of two main components, a system to render the physical shape output and a display for the spatially co-located 3D graphics. To view the spatially co-located 3D graphics, we utilize display arrangements well-known in AR: a stereoscopic spatial optical see-through display for single users (Figure 4-16a) and handheld video see-through displays for multiple users (Figure 4-16b).

The setup designed for single users renders 3D graphics on a stereoscopic display with a beam splitter, mounted on top of the shape display. When viewing the physical shape through the beam-splitter with spatially tracked shutter glasses, the graphics appear co-located. To explore collocated multi-user interactions, we also propose a version in which the 3D graphics are rendered on hand-held video see-through displays. While the graphics are not co-located in physical space, they are aligned with the video view of a camera mounted on the back of the tablet screen and therefore



(a) Single user setup: shape display and optical see-through display.

(b) Multi user setup: shape display and video see-through tablet display.

Figure 4-16: Sublimate setups.


provide the impression of co-location. As the display is handheld, it limits user interactions with the physical shape display to a single hand. While head-mounted displays could have overcome these limitations, we chose not to instrument the users head, as this would have limited face-to-face interactions. Another advantage of the handheld display is the built-in touchscreen, which provides an additional input modality for interacting with the content.

### **Shape Output with Optical See-Through Display**

Our first setup consists of a 2.5D shape display and a co-located semi-transparent 3D display. The shape display is based on a hardware setup similar to Relief [114], consisting of a table with 120 motorized pins extruding from the tabletop. The pins have a vertical travel of 100 mm and are arranged in a 12 x12 array with 38.6 mm spacing. The 3D graphics are rendered through a 120Hz 27" LCD screen with Nvidia 3D vision pro shutter glasses (60Hz per eye), which is mounted on top of a semi-transparent acrylic beam splitter. In addition to stereoscopic output, the user's head position is tracked by a Vicon motion capturing setup consisting of 10 cameras. This system creates a 425 \* 425 \* 100 mm space in which physical and graphical output are co-located for a single user. The shape display is controlled by a 2010 Mac Mini, which communicates with the application PC through OpenSoundControl (OSC). Applications and graphics rendering are running on a Dell Precision T3500 PC with a 2.53 GHz Xeon W3505, 8GB RAM and a Nvidia Quadro FX 4600 running Windows 7. All applications as well as the hardware control software are implemented in OpenFrameworks (OF). The system runs at 60fps.

### **Shape Output with Handheld Video See-Through Display**

To explore co-located multi-user interactions, we also built a version in which the co-located 3D graphics are displayed on handheld video see-through displays. We utilize 3rd generation iPads, which display a fullscreen video captured by its rear-mounted camera. A custom OF application tracks visual markers placed around the shape display using the Qualcomm Vuforia API. After computing the screen position



(a) Physically rendered surface with virtual control points. (b) Physical control handles with virtual surface.

Figure 4-17: NURBS Surface Modeling. The user can switch between rendering a) surface physically and control points virtually or b) control points physically and the surface virtually.

relative to the shape output, the video view is overlaid with adjusted 3D graphics. User input is synchronized between multiple iPads over wlan through OSC. The shape display is augmented with projection onto the object surface to enhance the object appearance and provide graphical feedback when viewing the shape without the iPad. The projector displays XGA graphics, which are rendered by a custom OF application running on a 2011 Macbook Air. The shape display is controlled by a 2010 Mac Mini, which communicates with the application computer through OSC. The system runs at 60fps.

#### 4.4.4 Applications

In order to highlight features of the sublimat system we created a number of example applications in different domains, such as computer aided design (CAD), geospatial data visualization and volumetric rendering of medical data. These different applications demonstrate the usefulness of state change of objects and interaction elements between physical and digital, augment shape displays to add resolution and scale, and provide augmented feedback to the user.

## 4.4.5 Single User Applications

### NURBS Surface Modeling

Manipulating 3D meshes with traditional 2D input devices, such as mice, is challenging, as their degrees of freedom are insufficient to control virtual objects directly. Gestural input has advantages due to more degrees of freedom, but lacks the material feedback of deforming real objects. We propose a basic application that combines physical control for mesh manipulation with an overlaid graphical view of the resulting surface (see Figure 4-17). The control points of a NURBS (Non-Uniform Rational Basis Spline) surface are represented by individual pins on the shape display. Moving the pins up and down affects the resulting surface, which is displayed through co-located 3D graphics. The control points are simultaneously highlighted through graphical feedback. The user can press a button to toggle the NURBS surface rendering from graphical to physical. In that case, the shape display renders the geometry of the modeled surface instead of the control points and the user can explore the physical shape with her hands. This application highlights the ability to dynamically sublimate control widgets, to allow for more precise control, or more degrees of freedom.

### Volumetric Medical Data Viewing

Volumetric data sets are rendered as 3D graphics which are spatially co-located with a physical shape in this application. The physical shape represents the bounds of the volume ray casting algorithm and can be reshaped by the user to create a non-planar cross section through the volume. This interaction is similar to Phoxel Space [159], but has the advantages of an actuated shape display, such as being able to save and load cross sections, as well as not being restricted to freehand manipulation. The cross section can be conveniently flattened and moved computationally, while the user can intervene at any time to modify its shape by hand. The location of 3D graphics is not restricted to the surface of the cross section, as volumetric data underneath or above the surface can be rendered to get a better understanding of the data set. This




Figure 4-18: Volumetric Medical Data Viewing. Users can modify cross-sections through the volume by physically deforming the object shape and switch between defined cross-sections through sublimation.

application shows how the system can quickly sublimate data to expose contextually meaningful areas; the user can, for example, annotate on a physically surface that has been revealed, more easily than in mid-air.

### **Virtual Wind Tunnel Simulation**

The virtual wind tunnel application renders different materials in their appropriate modality. While solid models are rendered on the physical shape display and can be touched and manipulated by the user, wind is rendered as spatially co-located 3D graphics. When the user deforms the physical model, a cellular fluid dynamics wind simulation updates accordingly. The wind flow around the model is visualized as transparent white lines floating in mid-air. To get a better view of the wind flow at a particular location, a tracked wand can be placed in the space around the model to disperse virtual smoke into the simulation, similar to the interaction proposed by Bryson et al. [17]. The virtual wind tunnel shows the advantages of augmenting shape displays with virtual graphics, and having bi-directional control of the output.




Figure 4-19: Virtual Wind Tunnel. The wind flow around a physical model formed by the user is visualized through graphics.

#### 4.4.6 Multi-User Applications

##### **Physical Terrain Model with Superimposed Virtual Information**

To test the multi-user setup, we developed an application for collaborative discussion of geospatial data. In this application scenario, the shape display renders physical terrain, while several tablet computers can be used to simultaneously interact with the physical surface. Seen through the camera of the tablets, we can expand the horizon of the physical map and display the terrain as it extends beyond the edges of the shape display. Pan and zoom controls on the tablets enable all users to change the region of interest. Moreover, users can display additional individual data overlays on their tablets. In our scenario, we provide a map showing radioactive contamination levels in Japan, as well as control parameters to drive a sea level simulation. The advantage of this configuration is that users can refer to the physical model during discussion with each other, while controlling a personal high-resolution view that allows them to switch between different perspectives of surrounding terrain or additional data layers.



Figure 4-20: Multi-user geospatial data exploration. Hand-held tablets augment the shape display by adding layers of data, and extending the active workspace.

## Chapter 5

# Remote Collaboration through Physical Telepresence

Much of our co-located collaborative work focuses on physical objects, surfaces, and spaces (paper documents, white boards, etc.) and interpersonal relationships that rely on physical interaction with other participants, such as a handshake. Computers have become important tools to aid such collaboration, but current user interfaces are often a source of disruption during face-to-face meetings [131]. In contrast, shape displays [131]. In contrast, shape displays can facilitate communication between multiple users, as they provide a shared spatial frame of reference, similar to the advantages of volumetric displays [56]. An example is shown in Figure 5-1a, where a shape display in a public setting is surrounded by visitors discussing the artifact with each other and touching the surface to collaboratively explore content.

Beyond co-located applications, we envision shape displays as tools that support real-time remote collaboration. As the world becomes increasingly connected, collaborative work is often distributed across multiple locations. A large number of commercial products are available to support video-mediated communication and shared digital documents. But many limitations for telepresence and tele-collaboration remain, as the affordances of the physical environment and presence of co-located collaborators are often missing in screen-based remote collaboration.

Researchers in distributed computer-supported collaborative work (CSCW) pro-



(a) Multiple users discussing a dynamic shape. (b) A local user discussing a digital car model with a remote user.

Figure 5-1: Co-located and remote collaboration through shape displays.

pose to improve video-mediated collaboration by situating it in the physical world through the use of shared media spaces [197] and projected Augmented Reality (AR) [223]. However, these interfaces still lack many physical aspects of collaboration, as remote participants are only visually present on the screen, which limits their ability to collaborate on physical objects. To overcome these challenges, telepresence robots have been proposed to embody remote participants for social presence and to manipulate the physical world from a distance [107]. A different approach to physical remote collaboration is presented by remote Tangible User Interfaces (TUIs), which focus on *synchronized distributed physical objects* [15]. These physical objects, or tokens, are synchronized with a remote counterpart to represent shared content, rather than embodying collaborators.

A goal of *Physical Telepresence* is to extend the physical embodiment of remote participants, which is common in telepresence robotics, and combine it with the physical embodiment of shared content, common in remote TUIs. An example of such a system is shown in Figure 5-1b, where the hands of a remote collaborator along with a shared digital model are materialized on a shape display.

This chapter presents different ways in which shape displays [153, 37] can physically embody remote people and objects to enable communication and collaboration. The flexibility of shape rendering also allows us to loosen the direct 1:1 link between remote shapes, both in terms of rendering users and data. After introducing

computer-mediation for physical telepresence, we propose how shape displays enable new forms of interactions that overcome the user’s physical limitations. We can, for example, replace a remote user’s arm with custom end effectors that are rendered on demand, as opposed to being limited to the static, prebuilt end effectors of a telepresence robot. We also propose interaction techniques that allow users to manipulate remote objects through gestures, mediating objects, tools, and direct touch. Finally, we report on a preliminary evaluation to explore potential use patterns in telemanipulation scenarios. The contributions in this chapter are:

- An exploration of physical telepresence for shared workspaces, using shape capture and rendering for rich, real-time physical embodiment and manipulation.
- Interaction techniques for manipulation of remote physical objects and shared, physically rendered, digital models.
- Interaction techniques that leverage physical computer-mediation to amplify user capabilities in remote operation.
- A technical software and hardware platform, which enables synchronous physical interaction for remote users, physical objects, and shape deformation.

## 5.1 Related Work in Distributed Collaboration

### 5.1.1 Video-Mediated Communication

Several early research projects investigate video-mediated collaboration for shared workspaces [198, 197, 79]. More recently, these techniques have been applied to applications, such as collaborative website development [30], remote board games [223] and family communications [225].

### 5.1.2 Mixed Reality

Collaborative tools can also leverage spatial interaction and mixed reality displays situated in the user’s environment. *TelePointer* by Mann [116] allows an expert to

point in a remote user’s real environment through a user-worn laser pointer, while more advanced Spatial AR techniques have been proposed for future collaborative workspaces, such as the *Office of the Future* by Raskar et al. [156] and the *Mirage Table* by Benko et al. Benko2012. Handheld AR can support mobile scenarios [186] whereas larger, situated displays have other benefits, such as immersion [54].

### 5.1.3 Telerobotics

A body of research focuses on representing remote people with telepresence robots (e.g., [206] [143]). The Personal Roving Presence concept [143] was an early exploration into tele-embodiment using, in this case, screen-based robots and flying blimps. Research has also explored the role of such devices in the workplace, and their influence on the sense of social presence [107]. Our approach of rendering captured geometry without semantic knowledge of user input, is quite different from common telepresence robots. Our interactions focus on leveraging the shape display for a rich, shared workspace, with less emphasis on mobility.

Another related domain of telerobotics is telemanipulation. Early master-slave manipulators used mechanical linkages, and now robotics arms and end effectors, to remotely handle hazardous materials or other objects at a distance [142]. A more specific example is that of telesurgery robots, which not only allow for remote operation, but give surgeons more precise control and remove tremors [188]. In contrast to prior work on telemanipulation, we focus on supporting collaborative work and the ability to switch between arbitrary end effectors on demand.

### 5.1.4 Tangible Remote Collaboration

TUIs can also be beneficial for remote collaboration. PsyBench [15] actuated a tangible object to mirror the movement of a remote physical token. Actuated Workbench [139] added more advanced control over multiple tokens, and represented remote user’s presence through projected digital shadows and haptic feedback. Researchers have also explored small tabletop robots for remote TUI collaboration [161, 163]. Video

portals have been used to play games over a distance, where physical pucks appearing from underneath the video create the illusion of a single shared object [125]. TUIs for remote collaboration often utilize actuated physical objects to represent content, rather than collaborators [15, 139, 161, 163]. However, as remote actors themselves are not physically embodied, object movement can result in a disconnected experience, since graphics can only partially simulate the presence of a co-located collaborator.

### 5.1.5 Shape Displays for Remote Collaboration

Since Project FEELEX [86], various form factors for 2.5D shape displays have been developed for co-located collaboration and interaction [114, 153, 37, 148]. Although Lumen has discussed the potential for remote presence and has demonstrated an example application [153], remote collaboration has been less of a focus in past shape display research [158].

## 5.2 Physical Telepresence

We introduce the concept of *Physical Telepresence* for shared workspaces, where capturing and rendering shapes has the potential to increase the sense of presence, expand the interaction bandwidth, and to extend human capabilities with computer-mediated interaction techniques.

Physical Telepresence is especially relevant for domains where physical presence, spatial information display and rich interaction is required. We primarily emphasize the potential in interactions with shared digital models of arbitrary shapes, linked physical objects, and manipulation of remote objects through direct gesture (see Figure 5-2).

We explore these interactions through three implemented systems:

- Bi-directional interaction through deformation on two linked shape displays (see Figure 5-17).




Figure 5-2: A user manipulating a sphere through his remote physical embodiment.

- Bi-directional interaction through shape capture and rendering: split-view on a single shape display (see Figure 5-3).
- Asymmetric teleoperation: A shape display linked to a video-mediated environment with shape capture (see Figure 5-16).

### 5.2.1 Example Scenario: Remote 3D design collaboration

Bill and Jean are two designers working on a new car body design at different locations using their respective Physical Telepresence systems. Bill sees Jean's face on a vertical screen and her hands physically rendered on a horizontal shape display; Jean has a similar view of Bill. As Jean moves her hands, Bill's shape display adapts to update the corresponding physical form of her hands.

Bill opens their shared car design model, which appears as a physical rendering. To bring the headlights in view, Jean gestures above it on her shape display, rotating and scaling the car, which updates the car model in real-time for Bill. When Bill deforms

the headlights with his hands, these changes in the shared model are propagated to Jean's display.

After discussing the changes, Bill moves the car to the side, and reaches for a headlamp that he 3D printed earlier. When placing it on the surface, a physical rendering appears on Jean's display. Jean rotates the rendering, which also rotates the 3D print on Bill's side. They make some digital annotations, which are projected over the physical part and the physical rendering. After reviewing necessary changes, they part ways, and Bill 3D prints the updated model.

### 5.2.2 Remote Physical Embodiment

The ability to provide relevant and adequate representations of remote participants and environments, which otherwise would have been co-located, is one of the grand challenges of telepresence systems. This embodiment has a great influence on realism, interaction capabilities, and sense of presence.

In a shape-display-mediated telepresence scenario, we identify four primary types of content to represent:

*Users.* Tracked remote and local users, with their representation transmitted to support interaction and telepresence.

*Environment.* The static or dynamic physical environment at each location, where the interaction takes place. We introduce a number of novel possibilities for computational control and actuation of the local and remote environment.

*Objects and Tools.* Tangible artifacts can be sensed in the environment, allowing them both to be manipulated by the system, or by a remote user through the shape display, in addition to direct manipulation by a local user.

*Shared Digital Models.* We can use a shape display to represent shared objects in multiple locations. Changes are updated in the computational model and propagated to remote locations. Shared Digital Models may have different viewports in different locations.




Figure 5-3: Top: Two linked tangible tokens at two locations, with synchronized position. Bottom: A tangible token rendered with the shape display at a remote location.

### 5.2.3 Representations: Physical and Virtual

Various techniques exist to capture the appearance, geometry and interactions from local participants and objects, which the system can use to form a representation that can be transmitted and materialized at a remote location.

*Physical Shape Rendering.* The physical shape of remote users, objects, and environment can be dynamically rendered through actuation on a shape display.

*Linked Tangible Objects.* Passive tangible objects, or tokens, placed on the surface of the table can represent remote objects. These passive objects can be manipulated through shape actuation [37]. This makes it possible to go beyond the limited degrees-of-freedom of the shape display by using external objects on its surface (see Figure 5-3, top, right).

*2D Graphics.* Different displays can be used to add graphics to the physical output. Previous work has explored integrating LED lights for color and illumination into shape display pins [153]. To overcome the low resolution of shape displays, high-resolution visual content can also be projected onto the shape rendering. This makes it possible to synchronize the 3D capture of geometry with a video stream that is

projection-mapped onto a dynamic, physical surface topology.

*Augmented Reality.* The previous chapter discussed how AR can add registered virtual content that extends beyond the range of the shape display [112]. This allows content that is floating in mid-air, has variable opacity, uses complex shapes, or has other properties that are not physically feasible or practical using current shape display technology.

*Complimentary Displays.* Other communication tools can be added to complement shape rendering, such as a vertical screen that give users face-to-face video. The video can, e.g., be aligned with the shape rendering to provide the illusion that the physical hands are coming out of the screen.

*Switching between Representations.* It may be advantageous to switch between these representations. A tangible token, e.g., could be represented by a remote tangible object, but if the remote user does not have enough tokens a physical rendering can be used instead (see Figure 5-3). Hands can be represented as physical renderings when manipulating an object, or as a virtual rendering when pointing or annotating (see Figure 5-9). These state transitions are in part inspired by our Sublimate concept [112].

## 5.3 Physically Manipulating Remote Objects

Prior work [37, 32] has demonstrated how shape displays can be used to manipulate physical objects, e.g., using translation and rotation. Here, we explore different control schemes to allow a user to manipulate remote objects. Using shape capture and display, users can reach through the network and pick up a remote physical object. These interaction techniques were developed through iterative prototyping with our inFORM system. Hundreds of visitors to our lab tried different techniques to manipulate objects remotely, and their comments and our observations lead us to develop and improve the following techniques.




Figure 5-4: Improvising with physical objects: Using a basket to scoop up remote balls.

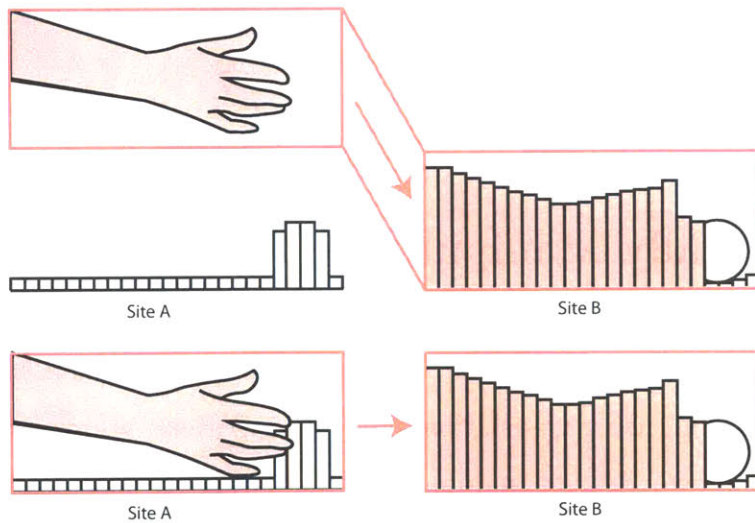


Figure 5-5: Transmitting gesture shapes directly above the surface, or from a separate mid-air zone above the shape output.

### 5.3.1 Gesture Control

*Direct Gesture Control* allows a user to interact directly with a remote tangible object through transmitted physical embodiment, which is rendered on the remote shape display. The rendered shape of the user directly applies a force to objects placed on the surface. For example, if the user forms a cup with the hands and raises them, this will be rendered and the pins could move a ball upwards. By observing the reaction of the physical object to their transmitted gestures, users can improvise to expressively manipulate a variety of objects. While our shape displays currently apply vertical forces, objects can still be translated laterally by tilting and sliding [37]. In addition to users' hands and arms, any object that is placed in the shape capture area can be transmitted and used to manipulate remote objects (see Figure 5-4). Because our shape display, based on inFORM [37], can only render shapes 0-100 mm in height, there is a question of how to render the remote environment, which most likely extends beyond 100 mm in height. We explored two mappings: Scaled and 1:1 with gesture zones. The scaled mapping takes all depth data and maps its height values to the shape display's maximum extent. The second mapping, 1:1 with gesture zones, takes some portion of the input space that is the same height as the maximum height travel and renders it on the shape display. This can be directly above the surface, or mid air, as shown in Figure 5-5. In the mid-air gesture zone mapping, users can touch the pin array without having their shape rendering automatically manipulate remote objects. When they want to manipulate an object, they move to the zone above it. In our asymmetric teleoperation scenario, we use a 1:1 mapping with the gesture zone of physical rendering starting right above the horizontal screen.

*Mediated Gesture Control* exploits that the remote object's pose is tracked and can be updated and moved to keep it synchronized with its underlying digital model [37]. *Physics-based Gestures* detect the user's collisions with the model, to update and move it, using a physics library. The updated model then causes the remote object to be physically moved. This is similar to the proxy-based approach in HoloDesk [67], but with actuated physical output. *Iconic Gestures* provide users with more abstract




Figure 5-6: The brush is equipped with a switch to detect when it is pushing against the side of a rendered object.

control. The user can pinch over the representation of the remote object to grab it, move their arm to another location, and open the pinch gesture to release it. The remote object is then actuated by the system to move to that location.

### 5.3.2 Interface Elements

Interface elements, such as virtual menus, can be projected around the rendering of a remote object to provide access to different operations. *Dynamically rendered physical affordances* [37], such as buttons or handles, that appear around the remote objects can be used for control. The user could press, push or pull such affordances to move the object.

### 5.3.3 Tangibles and Physical Tools

*Tangible Tokens* can be used to control a remote object. As the user moves the token, the model of the remote object is updated, and the remote object is moved to reflect the changed state. Two tangible tokens can be linked such that moving one causes the other to move, and vice versa, allowing for bi-directional control [139].

*Tools* allow users to manipulate remote objects by interacting with the local physical rendering. Tools can provide additional degrees of freedom (DOFs) of input, when interacting with the rendering. Our brush tool, for example, allows users to push re-




Figure 5-7: Pushing against the side of a shape-rendered objects with the brush tool moves it.

mote objects (see Figure 5-7). Figure 5-6 depicts a closeup of the switch and the bristles, which serve two purposes. First, they decrease the friction between the tool and the pins, which may get stuck when a lateral force is applied. Second, they smooth the haptic feedback resulting from the discrete, jerky motion when a physically rendered object is translated on a limited resolution shape display. To determine the direction in which to move the object, the brush tool's pose is tracked by an overhead camera, while a mechanical switch senses pressure.

## 5.4 Computer-Mediated Shape I/O

While shape capture and output literally adds a physical dimension to telecommunication, we find it interesting to go beyond representations that are symmetric in time and space. How can we create interfaces that enable remote participants to go beyond physically being there? The relaxation of 1:1 capture/output mapping opens up significant potential for new interaction techniques. Consider, for example a voice communication system with real-time translation of a speaker's voice to the listener's language (e.g. Microsoft Skype Translator). Here, the mediation emphasizes the relevant content and therefore simplifies the interaction between collaborators. Shape rendering allows us to apply similar concepts to physical remote collaboration. Another source of inspiration is the ability to augment the capabilities of the human body by adding or transforming it's physical shape, a concept explored through the




Figure 5-8: Scaling a users hand to interact with larger objects.

cyborg art of Stelarc [184]. More recently, Wu and Assada developed a wrist-mounted device that adds *Supernumerary Robotic (SR) Fingers* to a human hand [224]. These extra digits are synchronized with the hand movement of the wearer and support grasping, twisting and lifting actions to enhance human dexterity.

Inspired by these concepts, we explore mediating the physical shape rendering of the body, such as the transformation of physical form, altered representation, data filtering and replication, changing motion dynamics and time-domain manipulations.

#### 5.4.1 Transformation of Physical Form: Bending Body Limits

In our system, users can apply transformations to their activity in the remote environment, for example with scaling, translation, rotation, shearing, stretching and other distortions. Translation offsets geometry and can extend reach, with potential ergonomic benefits. Scaling can make a hand larger or smaller for manipulation of objects of varying sizes (see Figure 5-8). A small hand could avoid undesirable collisions in dense topologies, while an enlarged hand could carry multiple items. The transformations allow continuous real-time changes during the interaction, e.g., enabling smooth changes in size or position while holding an object. Examples of other transformations include replication or mirroring, e.g., to approach objects from multiple angles.

#### 5.4.2 Altered Representation: Becoming Something Else

With user and object tracking, there are benefits to switching representation for new capabilities. A system that captures geometry does not need to propagate all of it.




Figure 5-9: Copying and rendering multiple hands simultaneously (left) and switching to a virtual, non-physical, representation (right).




Figure 5-10: Replacing hands with a hook to reach or ramps to slide objects.

It may be useful to just send a user's hand and not the arm. As we are only limited by what the shape display can render, we can also morph into other tools that are optimal for the task, while controlled by the user. Examples include grippers, bowls, ramps, and claws — tools with specific properties that facilitate or constrain the interactions (see Figure 5-10 and 5-11). The tools could also be animated or semi-autonomously use the sensed geometry on the remote side to influence their behavior. Switching to a purely graphical representation to avoid collisions, is another example (see Figure 5-9).




Figure 5-11: The claw tool open and closed to enclose and move an object.

### 5.4.3 Filtering: Adding and Removing Motion

Signal processing can be applied to refine propagated motion, e.g., using smoothing or low-pass/high-pass filters. Such approaches are in use, e.g., in surgical robotics where hand tremors can be suppressed. Our system could also prevent fast motion or access to protected areas to avoid involuntary movements. In addition to reducing human noise, it may also alleviate system limitations, such as sampling resolution, speed, range and vibrations.

### 5.4.4 Motion Dynamics: Warping Time

Non-linear mapping of the propagated motion is interesting for many interactions. The properties of certain remote artifacts might, e.g., require slower or faster mapped motion, or require brief “freezing” or slow-down to emphasize an effect or make it legible. Such manipulations of time need, however, to be designed with great care, as they break the temporal link between the remote locations.

## 5.5 Prototype Applications

### 5.5.1 Telepresence Workspace

When discussing a physical design over distance, it is important for both parties to have an understanding of a shared model. We propose to render physical models on shape displays during remote collaboration meetings. The shape output is combined with video for viewing the upper body and face of remote collaborators. By aligning the tabletop shape display with the vertical screen, the two collaborators perceive a shared physical workspace, where the remote person can reach out of the vertical screen to physically point at a model. We support collaboration through shape rendering in several ways:

*Shared Digital Model.* The model is mirrored on the remote shape displays and provides a shared frame of reference.




Figure 5-12: Remote assistance: A remote expert helps a local user by supporting the workpiece with one hand and pointing with the other.

*Transmitted Physical Model.* When a user places a physical model onto the surface, its shape and texture is transmitted to the remote site.

*Physically-rendered Remote Collaborators.* When a user reaches out and places a hand on the table, its shape is physically rendered at the remote site. This conveys presence and enables remote pointing and manipulation. Figure 5-12 shows an example application, where a remote expert supports a local user.

### Unidirectional Shape Output

The shape output link described in the above scenarios can be uni- or bidirectional. An ideal bidirectional configuration, where both sites have a shape display, may not always be feasible due to size, cost and required infrastructure. A unidirectional link is still advantageous, as it allows remote participants to be more present with physically rendered body representations. It makes it possible to also capture and transmit physical objects to the remote site.

### 5.5.2 Collaborative 3D Modeling

A specific example of a telepresence workspace with a shared digital model is collaborative 3D modeling. In this case, users can deform the model, with any change




Figure 5-13: Collaborative 3D modeling of two linked landscape models.

reflected through the model on the connected site (see Figure 5-13). However, rendering the remote person’s hands may obstruct the changes to the model. Therefore, our application allows switching to rendering the model geometry without the user’s hands.

Another consideration is the potential conflict that may occur when two users try to simultaneously deform a model. We therefore use a turn-taking protocol for manipulation. When the system detects deformation, control is passed to the active user. The changes will be reflected on the remote connected model, while concurrent remote input is ignored. When the user stops modifying the model, the collaborator at the remote site may take control by starting a new deformation. Passing control between users can be applied to the whole model or to each pin individually.

### 5.5.3 Shared Presence: Connected Membranes

This application conveys the presence of remote participants through a shared digital membrane that uses two connected shape displays, each rendering the inverse shape of the other (see Figure 5-14). When a users pushes down a pin, the corresponding remote pin pops out. This is conceptually similar to a physical link between the pins. We implemented two modes:




Figure 5-14: Connected Membrane: When a user pushes into the membrane on one site, the inverse shape appears at the other.

*Pin screen.* Pins remain in their position after users push or pull them. This leaves behind a trace of past deformations, until it is erased.

*Elastic.* Pins always spring back to a neutral flat state when they are not being deformed. This mode conveys whether someone on the other end is pushing at that very moment.

## 5.6 Technical Implementation

### 5.6.1 Shape Output

Our proposed systems utilize 2.5D shape displays, which consist of arrays of linear actuators.

The *inFORM* platform [37] consists of  $30 \times 30$  individually actuated, plastic pins, built into a tabletop. The pins cover an area of  $381 \times 381$  mm with each pin measuring  $9.5 \times 9.5$  mm with 3.175 mm spacing between pins, and they can extend up to 100 mm in height. Actuation speed is 0.644 m/s and each pin can exert up to 1.08 Newtons.

*TRANSFORM* is a platform based on actuator modules with  $2 \times 12$  styrene pins on an area of  $305 \times 50.8$  mm, which extend up to 100 mm from the surface (see Figure




Figure 5-15: Side view of the TRANSFORM system with a single 12x2 actuation module.

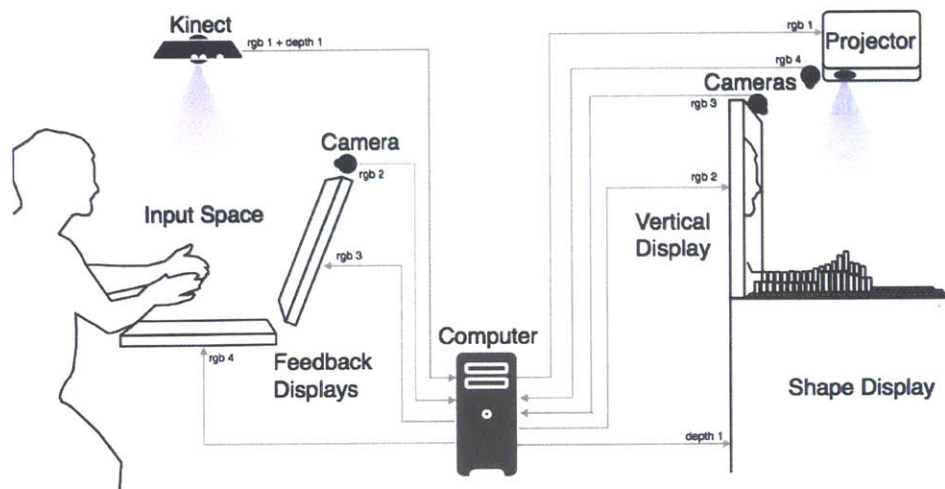


Figure 5-16: System diagram of the telemanipulation system using shape capture and remote shape display.




Figure 5-17: System diagram of connected workspaces with bidirectional shape capture and shape rendering.

5-15). The individual actuators and control boards are the same as used by inFORM. The modules can be seamlessly combined to form a larger shape display surface. We created two shape displays,  $16 \times 24$  pins each, covering an area of  $406.4 \times 610$  mm.

The shape display hardware uses custom Arduino boards that run a PID controller to sense and move the positions of 6 connected styrene pins through motorized slide potentiometers. Our applications are written in C++/OpenFrameworks and support 3D models (OBJ) and grayscale images as contents. The application renders a depth image, which is sent to the shape display over USB to RS485. Graphics from OpenGL are projected onto the white shape display pins through a calibrated ceiling-mounted projector (see [37] for details).

### 5.6.2 Sensing Shape Deformation Input

The pin positions that are reported from the shape display to the computer can be used to detect when a person deforms the surface. By comparing the shape image sent to the display and the position image received from the display, any deformation will be visible as a discrepancy. As the physical shape rendering contains some noise and the motors require a certain time to reach their target position, our detection algorithm thresholds the deformation image and compensates for the time delay. This

filtering is critical when sharing the depth information with a connected system, as the network delay and added noise can result in a feedback loop that results in false deformation detection.

### **5.6.3 Tracking People and Objects with a Depth Camera**

We track objects and gestures with a depth camera (first-generation Microsoft Kinect, 640×480@30 Hz). The 8-bit depth image is cropped, scaled and thresholded to fit to the size of the interaction area. The values are then normalized to match the range of the shape display. Our algorithm analyses the depth image to differentiate between the users' hands, shared tangible objects and arbitrary geometry. Hands are detected contours that touch the image boundary. After detecting the tip of the hand, its shape can be replaced by a tool. Shared tangible objects are colored spheres, which are tracked through their size and color.

### **5.6.4 Latency and Manipulating Objects**

The latency of our telemanipulation system, see Figure 5-16, is between 150 and 200 ms, depending on the amount and direction of travel for the pins. This is inline with latency used for other types of telemanipulation systems [115] and its impact is lessened by the lack of haptic feedback to the remote user. The latency is caused by a number of factors, such as the latency of the depth camera, depth image processing, communication bus, motor speed, and the camera and video latency for displaying to the remote user. Latency causes problems for keeping remote connected objects in sync, as shown in Figure 5-3. In order to adress this, we track the current position of the local object and move it at a controlled rate to the updated target position (dictated by the remote object).




Figure 5-18: Setup of the telemanipulation study. Left: The participant’s gestures are captured and two screens provide visual feedback. Right: Shape output moves the target ball.

## 5.7 Evaluation

Capturing the shape of a user through raw depth data and rendering it on a remote shape display is a new approach to physical telepresence.

Demonstrations at our lab of the teleoperation system to hundreds of users provided us with many interesting informal insights. Although the limited degrees of freedom of the system do not directly map to real-world operations such as grasping, lifting, and moving of objects, it does still provide meaningful interaction due to its real-world physics. We wished to examine this further in a controlled, qualitative study, to get a more nuanced picture of remote object manipulation strategies and interaction techniques.

In this preliminary evaluation we focus on teleoperation, since we believe that our techniques for remote physical manipulation introduce and enable significant new capabilities. We were particularly interested in the gestures that participants would use.

### 5.7.1 Apparatus

We used our telemanipulation setup (see Figure 5-16), where participants were in the same room as the shape display, but visually separated by a divider (see Figure




Figure 5-19: Participant’s view of the remote scene, with a top down 1:1 view on the horizontal screen, and a perspective view on the vertical.

5-18). The remote environment is shown with an orthographic 1:1 top-down view on a horizontal display, and a perspective view on a vertical display. An overhead, depth-sensing camera tracks the gestures above the horizontal display (see Figure 5-19).

### 5.7.2 Participants

We recruited 8 participants (2 female), ages 21–37 years old from our institution, who were compensated with 15 USD. All were daily computer users, 5 were monthly users of depth-sensing devices (e.g. Microsoft Kinect). Sessions lasted 45–60 min.

### 5.7.3 Task, Conditions

The three tasks required the participant to move a remote wooden sphere ( $\varnothing 4$  cm) to a target location ( $\varnothing 5$  cm). Users were not required to hold the ball in the target, only to touch (overlap greater than 50%) or cross it.

1. *2D docking*. The target location is projected on a random position on the remote surface, and 30 cm away from previous target for consistent task difficulty.
2. *3D docking*. The target location is randomized as in the earlier condition, but physically rendered at a height between 40–80 mm, requiring users to lift the sphere onto the target.
3. *Obstacle*. A variation of the 2D docking task, where an obstacle (water bottle,  $\varnothing$ 7 cm, 20 cm in height) was placed at the center of the surface. Study participants are to avoid knocking the bottle over.

There were four interaction technique conditions:

1. *Dominant hand*
2. *Dominant hand+*, which could be scaled, rotated or switched to virtual-only, by pushing buttons with the non-dominant hand.
3. *Two hands*.
4. *Claw* representation of dominant hand. Non-dominant hand could press a button to open and close the claw.

#### 5.7.4 Procedure

Participants had 5–10 minutes to familiarize themselves with the 4 different interaction techniques to manipulate the sphere remotely. Conditions were randomized and counterbalanced across participants. After each task, participants filled out a questionnaire to rate the task difficulty in each condition and briefly explain the rating. A post-test questionnaire had the participants provide additional comments. Video was recorded and reviewed to identify interaction patterns.

#### 5.7.5 Study Results and Observations

All users successfully completed all tasks. While some users noticed the system latency of 200 ms, they did not comment that it impeded them. However, we observed




Figure 5-20: Averages of users’ ratings of task difficulty for 4 different input conditions in each task. (7-point likert scale, 7 is best/easiest). Error Bars are standard error of the mean (SEM).

that all users adjusted their gesture speed when moving objects to compensate for latency effects. Because of our interest in emerging use patterns and a small sample size, we focus on reporting qualitative data and user feedback, as well as our observations of user interactions with the system.

The ratings (1–7, 7 is best/easiest) indicate that different techniques were preferred in different scenarios (Figure 5-20). Participants preferred using two hands in 2D docking, using the claw for 3D docking, and using one hand that could be transformed or the claw in the third, obstacle-avoiding task. The scores matched our observations and post-test interviews.

### Qualitative user feedback

Two hands for the *2D task* was intuitive and similar to how one would move a real ball. *“Moving the ball with two hands was a much more natural interaction for me”* Participants also used hand gestures similar to how they would in real life. *“I also felt slight stress on my hands while performing the task — for example, when I had to claw the ball, I could psychologically and physiologically feel the pressure of grabbing it, although it was virtual.”*

Using two hands was the most effective strategy in 2D, since participants could better enclose the ball, stop it, or move it. There were, however, other strategies: *“The biggest challenge was whether the cupping object (hand or claw) would sufficiently enclose the ball to prevent it from rolling away while moving it. The claw, scaled*

*hand, and two hands were all large enough to easily move the ball, but the scale control was the only mode that allowed for tuning the size."*

In the *3D docking task*, raising the ball without having it fall out of the remote hands, was a challenge for many participants. The claw tool was more popular here, as it constrained the ball well during 3D movement.

### **Cluttered spaces**

Real-world collaborative workspaces tend to have various physical objects on them, for instance cups, pens and paper. We were interested in how users would handle additional obstacles without haptic feedback. As expected, participants found this *obstacle task* the most challenging, while also the most realistic. *"Two hands were surprisingly frustrating in comparison to the prior tasks because they take up more space"*. But participants adapted to the challenge of avoiding obstacles very quickly by switching between a variety of different hand representations.

The claw's small footprint simplified navigation around objects. *"...perfect for this task as it enabled me to move the ball without having my arm get in the way. Rotating the view while using one hand also worked fairly well because it enabled me to adjust my trajectory to avoid the obstacle, although it was not as easy as the claw mode."*

Other participants preferred to turn off the hand rendering for obstacle avoidance. *"...most useful to go around the obstacles."* In general this technique was harder for people to discover. *"It took me a while to get a hang of the transparent feature, but it was good because it made my interaction faster, and give it a mouse like feel where I could hover and not click as when you use a mouse."* Participants also commented on picking up objects: *"I also found myself wanting a way to come from below an object – and didn't realize until much later that this is what the ghost/invisible mode enables that wasn't obvious how to use it to my benefit from the get-go."*

### **Observed use patterns**

We observed a number of techniques that the study participants employed to manipulate the remote object. To avoid bias, we had not given any prior instructions


on gestures. Here we outline the most prevalent for the single hand and two-handed conditions:

- *Push*: Hand (sideways like paddle), back of the hand, index finger, back of the arm, V-shape with fingers (thumb + index).
- *Hook*: Hook shape with hand (bring object closer).
- *Flicking*: Flick the object with an index finger.
- *Ghost*: Switch to virtual, move the hand underneath the object, then switch to shape rendering to pick it up.
- *Scoop*: Cup one or two hands, scoop up the object and move through the formed bowl. *"Two hands were easier to use when I had my palms close together. "*
- *Capture*: Approach with two opposing V-shaped hands to enclose the object, then move it around by moving the formed rectangle.

Our overall observation on the gestures emerging during the study was that participants adapted quickly to the degrees of freedom supported by the system and did not try to grasp the object. Instead, everyone handled the the remote sphere like a slippery object: pushing and flicking it sideways to translate it on the table surface, and scooping it with cupped hands to move in 3D. Some participants also experimented with other body parts and objects, like pushing with their arms, placing their head underneath the depth camera or experimenting with objects like pens and cups.

# Chapter 6

## Implementation



(a) inForm hardware setup.

(b) Transform hardware setup.

Figure 6-1: Hardware setups for inForm and Transform shape displays.

The research in this thesis was carried out on three different shape displays, which my collaborators and I built between 2009 and 2014: *Relief*, *inFORM*, *TRANSFORM*. *Relief* consists of 120 actuated pins arranged in a grid with 38mm spacing, extending from a circular tabletop [114]. *inForm* renders higher resolution shapes through 900 actuated pins arranged in a square grid with 13 mm spacing [37]. To achieve the denser pin spacing, the rods are connected to the actuators through flexible push-pull rods, which transmit bi-directional force through a bend. *TRANSFORM* consists of 1152 actuated styrene pins, embedded in a table surface and arranged in 3 individual grids

of 16x24 actuators, at 25.4 mm spacing between pins [81]. All these shape displays are augmented with additional sensing through a ceiling-mounted depth-sensing camera to capture objects and gestures.

All these systems required a combination of mechanical, electrical and software engineering. In this chapter, I will discuss the considerations that informed the hardware design, how each system was implemented, and how the design has evolved over multiple generations.

## 6.1 Actuation Platform

One of the most crucial technical choices is to pick an actuator for the shape display. As a large number of pins need to be stacked tightly, the actuator cost, footprint and complexity are important factors. Based on the design parameters discussed in Section 3.3, the following hardware criteria were defined for the mechanical design of the actuation platform:

1. *Actuation Speed:* Dynamic content should be rendered at a speed that is perceived as a continuous animation by the user. Speed is defined by two factors: the actuation speed and the bandwidth between the application and motor. Based on the criteria for stop-motion animation, 15 fps was defined as a minimum target refresh rate.
2. *Actuator Size:* The size of the individual actuators is critical for achieving a sufficient shape output resolution. In the cases where the actuator footprint is too large to fit directly under the surface, mechanical linkages can help to stack them efficiently. Another important consideration is the range of the actuator, or how far each pin can extend from the table surface.
3. *Actuation strength:* The actuators have to be strong enough to overcome the friction and weight of the pins and to be able to deform the top surface. They should however be weak enough to not hurt the user if they malfunction.
4. *Robustness:* The mechanical assembly and actuators must be able to withstand the user deforming the interface surface though pulling and pushing without

breaking. In addition, the control PCB's have to be protected from kickback current created by the motors when back-driving them.

5. *Power Requirements:* As we did not investigate mobile scenarios, power consumption is not a major consideration apart from the heat generated by the actuators.
6. *Parallelism:* To achieve a rendering and sensing speed of 15fps, the actuators and sensors of the interface should be able to work in parallel and independently of each other.
7. *Modularity:* To be able to handle the complexity of the number of actuators, a modular design is essential. Modularity speeds up the design stage and simplifies the building process, and allows for effective servicing of the interface. In addition, a modular design allows to reconfigure the physical form factor of the device into various actuator arrangements at a later stage.
8. *Cost:* To keep cost low, components like motors and mechanical linkages should be commonly available, and not require extensive modification for integration.
9. *Sensing Capabilities:* When a user presses into the shape to deform it, the system should accurately recognize the shape deformation and update the digital model. Previous CAD projects like *Sandscape* senses shape deformation through a ceiling mounted depth-sensing camera, which allows for a variety of materials and the capability to add external objects to the geometry. However, it has the drawback of the user's hands interfering with sensing while interacting. For this reason, an internal sensing approach was chosen for Relief, where the shape display actuators sense deformation and report them back to the system. This sensing is then augmented with external camer-based tracking to detect gestures, touch and objects.

### 6.1.1 Suitable Actuation Technologies

The actuation hardware and the control principles are similar to the technology employed in robotics and researchers can benefit from reviews for robotic actuators; an example is the overview and analysis by Hollerbach et al. [73]. Past projects have

utilized electric motors, electromagnetic actuators, shape memory alloy, piezos, and hydraulic and pneumatic actuators.

## Electric Motors

The most common actuation method for shape displays is electric motors, which perform a mechanical translation of rotary motion into linear motion. *FEELEX I*, *PINS*, and *ShapeClips* combine motors with lead screws, while *Surface Display* and *FEELEX II* use geared servo motors with a crank mechanism.

*Motorized slide potentiometers* are commonly employed in motorized mixing boards. These actuators are based on linear potentiometers equipped with a motor and an internal belt to move the potentiometer lever to a target position. Their advantage is that they are commercially available, relatively inexpensive, and robust, as their lack of a gearbox allows backdriving without them braking. A disadvantage is that they are only able to drive small loads with up to 1N force.

Novel linear DC servomotors (for example the *Faulhaber Quickshaft LM127*) have the advantage of high speed and actuation strength with a small footprint. While the price and availability of such actuators limits their applicability to shape displays at this time, they could overcome some of the limitations of current electrical drivetrains in the future [25].

The footprint of the electromotor is often too large to fit underneath a single pin. Therefore, some displays use Bowden linkages to be able to stack motors on top of each other. Examples are *inFORM*, *Transform*, and *Emerge*. As the linkages and stacked actuators lead to a taller overall structure for the shape display they result in a large overall actuation assembly. In addition, the length of the straight pin has to be taken into account, which adds to the height. A novel approach to overcome this limitation is to utilize flexible, rolled up pins, similar to tape measures, such as proposed for the *KineReels* system by Takei et al. [195].

## **Electromagnetic Actuation**

Electromagnetic actuators like solenoid or voice-coil actuators have the advantage of a potentially smaller package that fits underneath individual pins. As solenoids only allow for an on/off control mechanism, so the each pin can only be in one of two positions - an example is the *Shade Pixel* project. Voice coil actuators have the advantage of finer granularity of control, but are more expensive and complex to control. *ForceForm* uses a similar approach based on electromagnets, while *BubbleWrap* uses an arrangement similar to a voice coil, with movable solenoids embedded in a soft fabric.

## **Shape Memory Alloys**

Displays that utilize shape memory alloys are *Pop Up!* and *LUMEN*. While SMA allows for actuators with a small footprint, it's thermal properties may require active cooling to allow for precise control and speed up the response time [127].

## **Piezoelectric Actuation**

While Piezo actuators do not provide sufficient displacement for shape displays, this technology is commonly used for tactile displays. But novel piezo actuators that utilize friction clutches for moving a longer pin may be a suitable future technology.

## **Fluid: Pneumatic and Hydraulic**

*Hyposurface* [53] and *Gemotion Screen* [132] are based on pneumatic actuators, while *Digital Clay* investigates the use of hydraulics [228]. An advantage of fluid systems is that the actuation system can be distributed: pumps and reservoirs can be placed in a different location from the individual pin actuators without requiring a complex mechanical linkage between them. This enables the design of a system without bulky and noisy motors underneath the surface. In addition, this technology is well-suited for soft, transparent actuators, such as button overlays for multitouch displays [62, 21]. A disadvantage of this approach is the system complexity and potential fluid leakage.




Figure 6-2: RSA0N11M Motor-driven master type slide potentiometer by ALPS [2].

## Summary

To balance fast actuation speed, build in sensing, and sufficient stroke, while allowing for user interaction like pushing and pulling without breaking, we decided to use the RSA0N11M motorized slide potentiometers by *ALPS* (see 6-2). So goes the engineering part of the story – in reality the fact that we got a good deal on a surplus of these actuators on eBay for USD 4.00 per piece also played a significant role. The comparably low cost, the above performance characteristics and the fact that they could be easily controlled via Arduino microcontroller through attached motor shields outweighed the disadvantages like low actuation strength and a non-ideal footprint. When designing our next generation shape display, *inFORM*, these actuators had performed admirably, with a  $< 2\%$  failure rate during two years of lab use and frequent demonstrations. We therefore resorted to using similar actuators for all our subsequent shape displays.

## 6.2 Table Hardware and Setup

### 6.2.1 Relief

When designing *Relief* in 2009, its main inspiration came from previous physical CAD interfaces developed at the Tangible Media Group, such as *Illuminating Clay* [147]. This background and the application domain of landscape design informed

the initial size, tabletop form factor, continuous shape display surface and hardware functionality of the interface. The input surface and sensing capabilities of the system evolved later on, as new applications were developed.

The round table form factor of *Relief* was chosen so multiple participants would be able to approach and interact with it from different directions. As the table hardware needed to be simple to fabricate and the environmental impact should be kept at a minimum, we used plywood for the table frame. This choice of material allowed for rapid prototyping by laser cutting and machining parts with a Shopbot CNC router. The control electronics consist of 32 Arduino Duemillanove boards with attached Ladyada motor driver shields, each powering 4 actuators, which are mounted into individual plywood module housings. These modules can be rearranged into different physical layouts, like a square, rectangular or circular footprint. At the end of the project's lifetime, the table frame can be disassembled, and the electrical components removed and reused for future projects.

The surface of the interface was meant to form a continuous landscape model; therefore, the individual pins were covered in stretchy white Velcro fabric. To attach the fabric surface to the actuators, we used several different methods: At first, we resorted to a magnetic connection, with a neodymium magnet attached to each pin and a spherical neodymium magnet on top of the fabric. The magnets had the advantage that they allowed the fabric to slide when stretched and therefore allowing the fabric to comply better with the rendered shape. They also provided small handles to grasp and pull at if the user wanted to move a pin up. However, they interfered with the projection and aesthetics of the interface, and it was tedious to remove and reattach the magnets each time the interface required servicing. Later, we removed the magnets and resorted to glued 3M velcro connectors to attach the fabric to the pins.

*Additional Tangible Input:* A large ring controller surrounding the display was used to add an additional input dimension besides pulling and pushing. By turning the ring users could, switch between models, zoom in and out, or rotate geometry. The rotation of the ring controller was sensed through an optical mouse mounted

upside down.

## 6.2.2 inFORM

The design of the second generation of shape display, inFORM, was based on the previous experiences of building and using Relief, which had been in continuous operation for 2 years at that point. While the motorized potentiometers had performed well, the main goal was to improve the resolution and to change the form factor of the interface surface into a square shape composed of a grid of square plastic top pins.

The system uses  $30 \times 30$  motorized white polystyrene pins, in a  $381 \times 381$  mm area. The pins have a  $9.525 \text{ mm}^2$  footprint, with 3.175 mm inter-pin spacing, and can extend up to 100 mm from the surface. *Sullivan Gold-N-Rods* Push-Pull rods are used to link each pin with an actuator, to enable a denser pin arrangement than the actuator footprint would allow, giving the system a height of 1100 mm. The linkage, a nylon rod inside a plastic housing, transmits bi-directional force from a motorized slide potentiometer *ALPS RSA0N11M9A07*, through a bend. Six slide potentiometers are mounted onto a custom-designed PCB, powered by an Atmel ATmega 2560, and TB6612FNGCT-ND motor drivers. The linear positions are read by the 10-bit A/D converters on the microcontroller, and allow for user input, in addition to servoing their position using PID control.

150 boards are arranged in 15 rows of vertical panels, each with  $5 \times 2$  boards. The boards communicate with a PC over five RS485 buses bridged to USB. The system has a 60 Hz refresh rate, determined by the 115200 bps RS485 bus speed, the 8 byte control message, and 30 boards on each RS485 bus.

For each pin, we can update both position and PID terms to provide haptic feedback and variable stiffness, for example, to create haptic detents or buttons that are harder to press. This control also allows us to limit power consumption or to switch off individual pins to avoid burning out the motors.

Pin height and stiffness is represented in software as an 8-bit height map, which can be produced by OpenGL shaders, through different shape primitive classes, or by directly writing values into the height map. The height map is then mapped to




Figure 6-3: The inFORM system actuates and detects shape change with 900 mechanical actuators, while user interaction and objects are tracked with an overhead depth camera. A projector provides visual feedback.

individual pins and sent to the microcontrollers. Similarly, the feedback from the 900 pins is received over the RS485 bus, and used to track user's surface deformations as input.

Each pin can exert a force of 1.08 Newtons (equivalent to 100 g weight), which was measured using a precision digital scale. The effective average upwards and downwards speeds (0.644 m/s and 0.968 m/s, respectively) were measured using a high speed camera.

In theory, the system's 900 pins could consume up to 2700 W due to a peak 3 W power consumption per actuator. In practice, however, we measure the effect to

approximately 700 W when in motion. Due to friction in the linkages and actuators, the pins maintain their position when unpowered. Therefore, the power consumption to maintain a shape is less than 300 W (mainly due to the 83 % efficiency of our 12 V power supplies). Heat dissipation remains a critical design criteria for the actuation assembly and we use two rows of five 120 mm fans to cool the actuators.

### 6.2.3 TRANSFORM

*TRANSFORM* was conceived as a study of dynamic furniture at the *Lexus Design Amazing* exhibit during the *Milano Design Week* 2014. As it was running continuously for 8 hours a day for a week, the main focus of the mechanical design was to minimize friction, make the assembly more robust, and allow for faster swapping of broken actuation components.

We also changed the mounting structure and design of the top pins to minimize the gap between pins.

### 6.2.4 inFORM for “Tools” exhibit

In 2014, we built a new version of inFORM that was part of the 6 month exhibit “Tools: Extending our Reach” at the *Cooper Hewitt National Design Museum*. Through this shape display, museum visitors could experience four interactive applications: physical telepresence, mathematical formulas, cityscapes, and a wave simulation.

The main focus of this project was to create a shape rendering mechanism robust enough for daily use in a museum setting. As the moving parts of the shape rendering mechanism should be protected against damage and spills over the 6 month period of the exhibit, the shape display was enclosed with plexiglass sides (see Figure 6-4). Instead of interacting directly with the shape display, visitors could control the applications through a kiosk with a touch screen and Kinect camera. Another safety measure was to switch off the motors when no visitor was nearby. If the system detected the proximity of a visitor through the Kinect camera, it would turn on the shape rendering and switch to an interactive mode.




Figure 6-4: The inFORM setup created for the “Tools” exhibit.

### 6.2.5 Rendering

Figure 6-5 depicts the basic rendering pipeline. The depth map and graphics are rendered in OpenGL using openFrameworks [136] and sent to the shape display and projector. The projector’s world coordinates were determined after calibration with the Kinect color camera, while distortion coefficients and intrinsic parameters were recovered using ProCamCalib [4]. The projector’s  $1400 \times 1050$  pixel output is then warped to correct for the parameters and projected over an  $87 \times 66$  cm area (in the case of inFORM).

### 6.2.6 User and Object Tracking

Our current system uses an overhead depth camera to track users’ hands and surface objects, as shown in Figure 6-3. A Microsoft Kinect with a  $640 \times 480$  pixel depth sensor is mounted 1200 mm above the surface and calibrated for extrinsic and intrinsic camera parameters. We combine a static background image of the table surface with the shape display real-time height map to form a dynamic background image

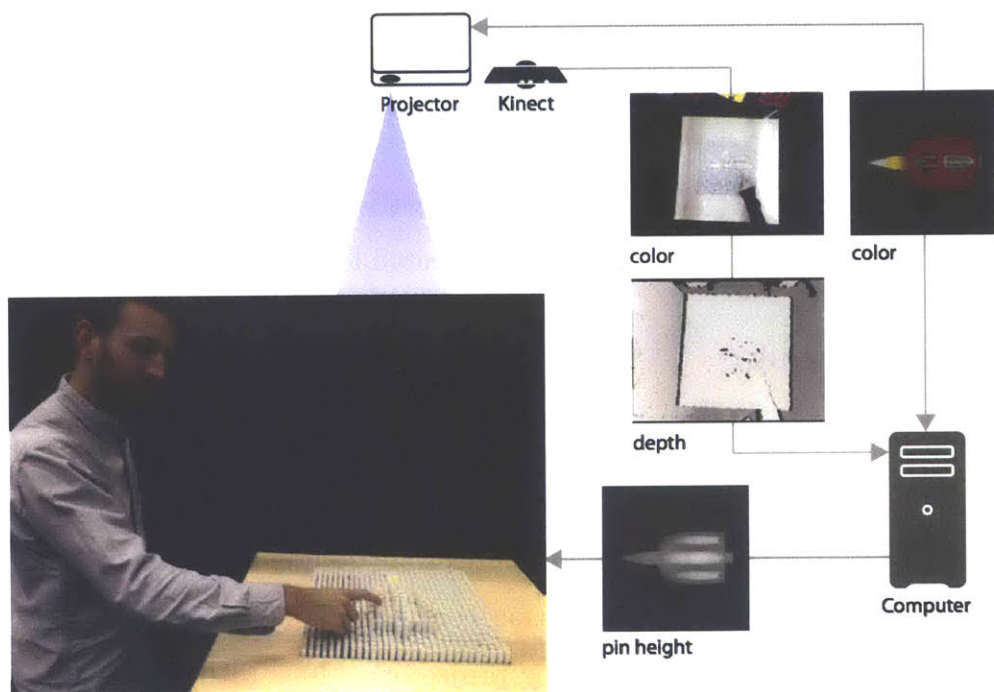


Figure 6-5: Basic pipeline for rendering shapes and graphics and tracking user input with a depth camera.

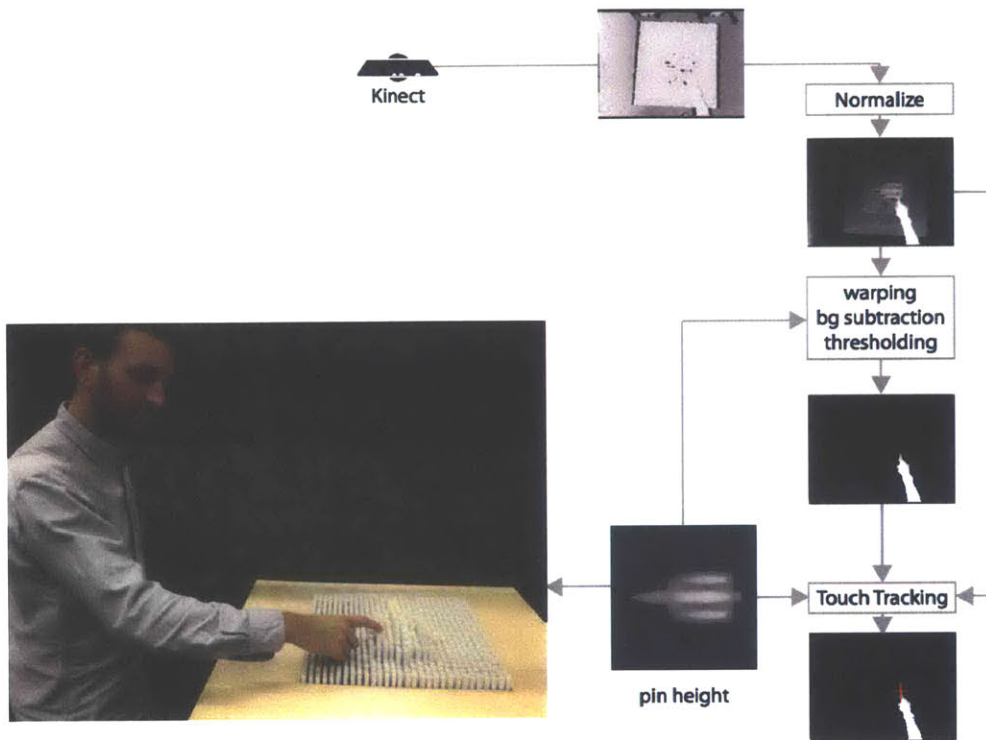


Figure 6-6: Tracking pipeline to detect the users gestures and touch.

that is used for subtraction and segmentation. This height map is scaled and a homography is applied to warp the image into camera space. We find all pixels above the computed dynamic background model, and threshold the image. For hand and fingertip tracking, we use OpenCV to compute contours of the threshold image, followed by convex hull and convexity defects. The 3D finger tip coordinates are transformed to surface space, and an arbitrary number of finger tip contacts can be tracked and detected to trigger touch events (see Figure 6-6). For object tracking, we currently rely on color tracking in the HSV color space (see Figure 6-7). However, other approaches to object tracking would be straightforward to implement. Objects and finger positions are tracked at 2 mm resolution in the 2D surface plane, and at 10 mm in height. The touch tracking can be combined with input from the shape display actuators to distinguish button presses from light touch.

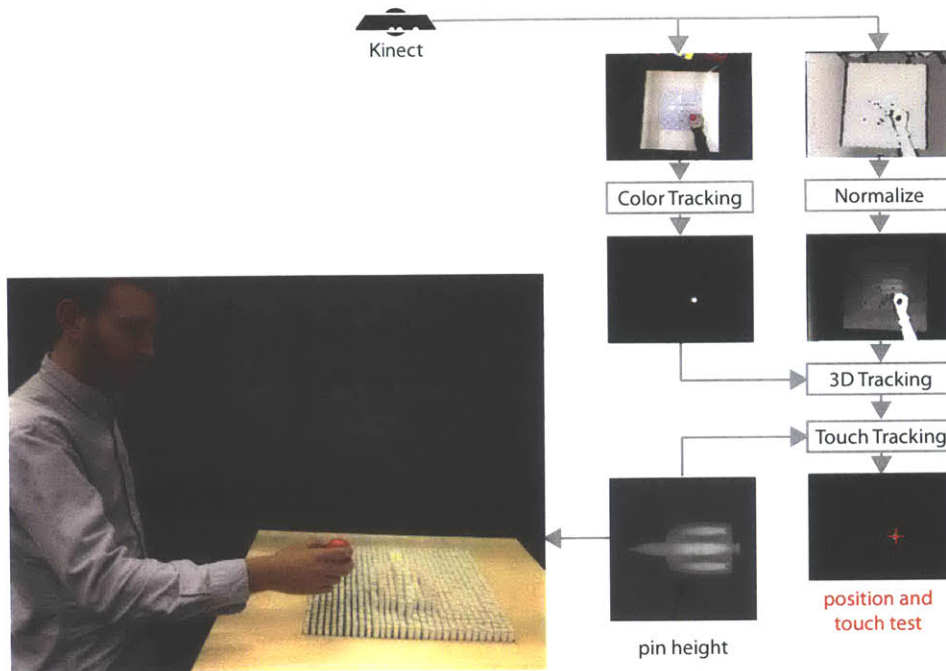


Figure 6-7: Tracking pipeline to detect object positions.

### 6.2.7 Future Touch Detection

As an alternative to the limited depth camera touch tracking, it may be interesting to embed touch sensors, e.g., capacitive or optical, directly in the pins. More sophisticated object tracking, for example, using multiple cameras, markers, or sensing in the electromagnetic or RF domain, could be used to address user occlusion, object identification, and color interference with projector.

# Chapter 7

## Discussion

In the process of building shape displays over the span of six years, we formed a set of general guidelines for application development. We also built tools to aid researchers and artists when using shape displays as a storytelling platform and to animate content. We think that these guidelines and tools are useful beyond our own systems and discuss the lessons we learned so far.

### 7.1 Design Guidelines

#### 7.1.1 Interaction and Rendering beyond the Surface

In our systems, we utilize direct touch to modify spatially relevant parameters, and mid-air gestures to control more global, abstract and view-oriented parameters. As spatial interaction moves beyond direct touch, not all information is best represented physically - as researchers we should consider allowing users to switch the representation of objects from physical to co-located 3D graphics.

An important constraint is the bounds of a shape display's actuator range. 3D information might fit in the surface plane, but will often extend beyond the height limits of the actuator. Should the object be cropped, squashed, or scaled? While some of these transformations can be decided by the application designer, giving the user the tools to adjust the scale of data interactively is a powerful way to adapt them

to their needs and hardware. The *ShapePort* interactions described in Section 4.3 can provide such a capability. Additionally, spatial 3D graphics, such as the *Sublimate* system can render visuals that exceed the physical boundary of shape output (see Section 4.4).

### 7.1.2 Improvising through Inter-material Interaction

Because shape displays render physical surfaces, objects can be placed on them and the rules of physics apply to their interaction. Designers can leverage this effect to provide users with tools to understand data through their intuition of naïve physics. For example, a ball placed on the surface graph of a function will roll towards a local minimum, which can aid in understanding the gradients of the function. We observed that users often started to play and improvise with found physical objects, such as their cellphones. Regardless of whether an application was designed for such interactions or not, they would place objects on top of the moving rendered shapes to see how they interact or they would use them as physical tools to press down multiple pins at once. Application designers should anticipate and embrace such improvisation with found objects.

### 7.1.3 Animation Principles

In our user studies and by observing everyday interactions with our shape displays, we noticed that users can experience shape change as jarring [114]. A common example is when switching between applications, or when the system loads a new object that suddenly appears on the table. As we intend for our shape transformations to delight rather than startle users, we explore a number of different animation techniques. Some of these are informed by traditional principles in animation [105], and by theories of movement in HCI, like Vaughan’s treatise on understanding movement [213] and Young’s adaptation [165].

*Speed:* The first and simplest method we experimented with was to change the speed of shape change. Slowing down shape transitions makes them less jarring, but

speed does not need to be a single global parameter. We noticed that fast movement is more jarring when it appears closer to the users hand. Therefore, the speed of shape change of an object can be programmed to progressively adapt to the distance to the user.

*Motion Quality:* Instead of moving at linear velocity, object movement is defined by splines that “Slow In and Slow Out” velocity at their start and end position to mimic the behavior of physical objects accelerating and decelerating [105]. Giving objects a more natural motion makes them less uncanny, as it allows users to anticipate their movement. Other parameters to achieve more natural motion are added noise and secondary motion [140].

*Scale:* We noticed that small shapes are less jarring than large shapes moving at the same speed. We can use this observation to scale objects to a smaller size, move them and then scale them up to their final size again. Scaling objects was also a successful strategy when loading new objects on the display.

*Direction:* The direction that the shape moves makes a difference to the user: we therefore experimented with directional movement on *inFORM*, where objects appear to move in from the side of the table, rather than suddenly materializing in the center.

*Graphics:* Impending motion can be visualized through projected or spatial 3D graphics. A system like *Sublimate* can render a graphical preview of physical shapes before they appear. Such visualizations can apply metaphors from science fiction like the particles of light that appear around objects as they are materializing - like the well-known example of the *Star Trek* transporter beam [103].

*Control and Interactivity:* When users have control over the shape change, they usually enjoy rapid motion rather than find it disturbing. An example we often observed is the user interaction with the fluid simulation on the *TRANSFORM* table. It produces large, rapidly moving waves when users gesture around it with fast movements, but quiets down when users stand still. Creating a “big splash” on the table would often bring a smile to the users face, which indicates that when the causality of how the motion relates to the body is clear, it leads to a delightful experience.

## 7.2 Tools for Artists and Researchers

### 7.2.1 A Platform for Design Fiction

HCI has always been fueled by research visions that describe the potential of future technologies, take the example of Vannevar Bush’s 1945 article “As We May Think” [18] or Ivan Sutherland’s description of “The Ultimate Display” in 1965 [191], which inspired decades of subsequent research. Beyond these visions that came from an academic context, popular science fiction has often provided inspiration for technology. The shape display in the movie *X-Men* [181] inspired the development of shape displays like the *Xenotran XenoVision Mark III* system [179]. Within the HCI community, video prototypes have become into an important tool to overcome hardware limitations and effectively communicate a user experience, as argued by Tognazzini in the discussion of the *Starfire* video [202]. This approach has been further developed through practices like *Design Fiction* [196], *Diegetic Prototypes* [97], and *Speculative Design* [28] to ignite discussions on future implications of technology from a design perspective.

These practices are highly relevant in the context of research on shape displays, since current technology is still far from an ideal shape-changing material. The modular robots depicted in *Claytronics*, Radical Atoms in *Perfect Red*, and hypothetical nano-materials in *Nano Futures* are examples of videos from the HCI community that envision hypothetical interactions with future materials. But such video scenarios are only successful if they are plausible, and therefore require considerable time and effort to produce. In our own research, we have created video prototypes like *Contour* to simulate future shape displays. Through stop-motion animation with non-functional props, we simulated an actuated, high resolution 2.5d surface and user interactions with it. Using physical props instead of CGI not only aided in filming, but also informed our process significantly and produced new research ideas. But modifying these props by hand for each frame of the stop motion animation slowed down our process and we realized that we wanted a more flexible physical tool to iterate and quickly discard user interactions while we developed and filmed them. Later, as we




Figure 7-1: Still image of the video prototype *TRANSFORM as Dynamic and Adaptive Furniture* [216].

built our own shape display hardware, we started to utilize these systems as such storytelling tools. An example is the video prototype of *TRANSFORM as Dynamic and Adaptive Furniture*, which proposes how shape-changing surfaces may become a part of everyday life in the future [216]. In the video, actors interact with physical props while pre-scripted animations are played back on the shape display (see Figure 7-1). These animations were faster to produce and easier to modify than our prior stop-motion techniques, and they allowed for a more convincing end result than post-production 3D renderings would have. We think that future shape displays, with more degrees of freedom, can push such a storytelling approach even further. To make full use of this platform, animators can utilize various tools for programming shape change, discussed in the next section.

### 7.2.2 New Brushes for a Transforming Canvas

When developing applications and telling stories through shape displays, we need expressive tools to design how shapes transform. Metaphorically speaking, these tools are the digital brushes that allow us to paint shape motion onto the canvas of the

shape display. The ideal tool depends on the type of application: while some applications only incorporate minimal motion design, others rely on an extensive and careful choreography of the pin movement. An example of minimal pre-designed motion is the *Recompose* application, where only the speed at which the initial models load is defined programmatically - all other motion is controlled by the users gestures and touch during runtime. On the other end of the spectrum, *TRANSFORM* requires extensive motion design, as it combines pre-defined motion with procedural animations. While in the former case, programmatic tweaking of a single motion parameter suffices, the later application benefits from more sophisticated tools to design the speed at which animated objects appear and transform. The tools we developed for our projects range from software for blending keyframes, to 3D animation software plugins for controlling the shape display, tools for recording and playing back user motion, and adding secondary motion through simulated materials. We believe that these tools not only reflect our own practice, but that their discussion points towards a rich set of future brushes that artists and HCI practitioners will develop for shape-changing interfaces.

## Software Tools

*Blending Keyframes:* An important software tool developed for *TRANSFORM* was to define a class structure where every object on the shape display automatically possesses a set of animation parameters and can be blended with other objects. That means that any content, whether it is a 3D model, image, procedurally generated shape, video, or camera input, has a start and end position with an alpha value and an animation curve. Blending between these diverse objects is enabled by blending the depth textures that the system generates from them before sending them to the display. With this tool, designers can animate a variety of content with little effort: content can for instance dramatically rise while from the table surface while slowly rotating when loaded, or objects blend into other content when the user interacts with them. Besides programmatically controlling such transitions, we also implemented a timeline GUI to control animations and add interactivity. For this purpose, we have

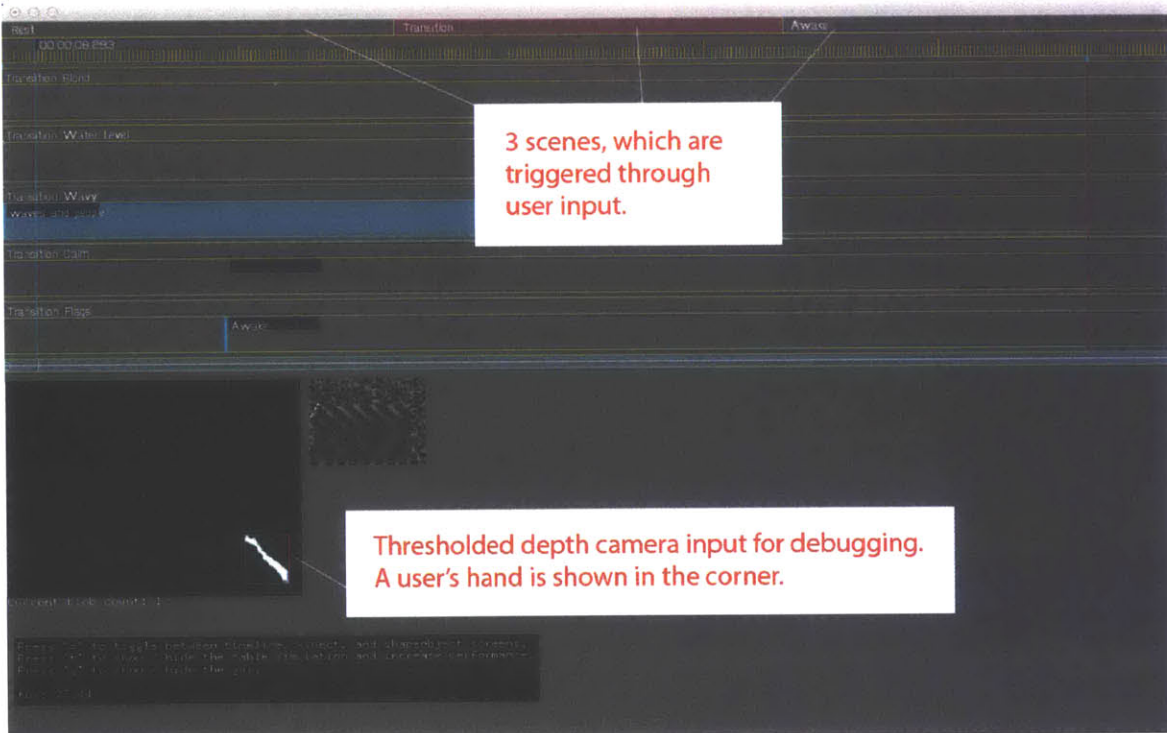


Figure 7-2: TRANSFORM timeline GUI for triggering and blending animations.

adopted the `ofxTimeline` [47] add-on for `openFrameworks` [136] to our software platform, which provides an interface for keyframe animation, hardware input triggers, and for loading and managing content (see Figure 7-2).

*Pre-rendered Animations and 3D Animation Software:* Besides blending content, the system can load and play back pre-rendered animations from video files containing depth and color textures that are exported from animation software. To speed up development, Philipp Schoessler developed a *3ds Max* [221] scene of *TRANSFORM* to pre-visualize the shape display transformations for the animator. However, viewing an animation on a computer screen is very different from experiencing it come to life on the shape display, and the extra step of exporting, copying and loading a video file to test animations on the actual hardware added a tedious hurdle to the process. For this reason, Philip also developed a plugin for *Autodesk 3ds Max 2014* that acted as a realtime networking bridge to the shape display hardware (see Figure 7-3). Using this software, models and animations can be played back on the shape display during the creation process.




Figure 7-3: *3ds Max* networking plugin to play back animations on *TRANSFORM*.

## Embodied Animation

The software tools are useful for development, but fall short from our goal of using the shape display itself as a tool for expression. To achieve this, motion design for the interface should be informed by the user directly through the interface. This can be directly, through deforming the interface surface, through objects, or through mid-air gestures.

*Record and Playback of Hand Gestures and Materials:* By recording how designers deform the interface surface, they can create animations similar to the kinetic memory functionality of the *Topobo* system by Raffle et al. [155], where their physical pin movement is repeated after they animate it. Besides pressing into the pins with bare hands, animators can also use tools and objects to press into the table. This allows them to for instance modify many pins at once or create types of animations that are not possible with bare hands. Besides deformation, free-hand gestures can be utilized to drive animation. Using the physical telepresence setup to capturing shape and replicate them on the shape display, we created a tool to record and play back user motions. These shapes are often free-hand or body gestures, but we experimented



Figure 7-4: An animator deforming a fabric surface to inform shape change [38].

with physical props and materials as well. Figure 7-4 depicts an example of such a material interface, where a fabric is stretched over a frame and deformed from underneath. Sean Follmer utilized this setup for a study in which participants were asked to express emotions by deforming the fabric surface [38]. During the recording session, the participants' movements are rendered in real-time on a connected shape display to provide them with physical feedback. The movement is stored as a video file and can be played back and looped on the shape display later on.

*Secondary Motion Through Simulated Materials:* While real materials give animators expressivity during the recording process, simulated materials can add rich behavior through secondary motions when played back on the shape display. Examples are rigid-body dynamics, simulated fabrics or fluid systems. Figure 7-5 shows a water simulation reacting to mid-air gestures above the display. This system could be extended with further simulated materials to create a rich library of behaviors for the animator to use. After defining the primary motion of an object, designers could select such a simulated material for secondary motion and tweak its behavior to create a richer animation.



Figure 7-5: A wave animation is driven by the user's mid-air gestures.

## Future Work

The methods described above are currently the most comprehensive set of tools for generating and modifying animations on shape displays. These tools and guidelines were developed for our custom tabletop shape displays, but they could be adapted to very different form factors, like mobile devices, walls for interactive architecture or an animated theatre stage for live performances.

In the future, the control that these tools give to animators for modifying recorded shape change later on could improve. Right now, animators can modify the playback speed and loop movement, but do not have fine-grained control of physically adjusting the playback animation curves later on. We envision that such controls could be linked to the animators gestures: they would first define a start and end position for an object on the table, then play out the motion curve for the object. Physical objects on top of the shape display could act as handles for the animator. A future system could also incorporate a more sophisticated physics engine to simulate various material behaviors

of objects rendered on the shape display. A repository of materials, behaviors and animations could be shared between multiple animators online and applied to the behavior of applications, similar to style sheets.

Apart from recorded motion, the real-time functionality of our *Physical Telepresence* workspaces could be extended to support “Wizard of Oz” experiments. A designer could play out how an interface reacts to user input to quickly prototype and test interface transformations.

While the tools described in this section have been developed for artists and HCI researchers, they could also enable end-users to personalize how their products react. Imagine for instance a future mobile phone that transforms its shape according to personal preferences, or in a way that resembles the gesture of a person close to us.

### 7.3 Technical Improvements

Many of the limitations of shape displays are related to the current shape display hardware, such as the degrees of freedom (DOF) and resolution.

*2.5D rendering* using physical pins limits our system’s DOFs to rendering reliefs and prevents overhangs. This is especially important for telemanipulation, since it only allows application of vertical forces to objects on top of the pins. The system can thus not grasp objects [24], only lift or translate them by tilting and sliding. Robotic grippers could be combined with shape displays to provide more dexterous manipulation. Shape displays with more DOFs of output can also be explored to allow lateral forces with more intricate control. Latency, as discussed, is another parameter that limits manipulation, and further improvements in hardware and interaction techniques can address this.

Similarly, shape displays have limited degrees of freedom of input; users can only push down and pull up on the pins. In addition, each pin is connected to the shape display, which means that users cannot pick up or physically move parts. This limits the degrees of freedom and the expressiveness of input. We address this to some extent through tangible objects on top of the display, but future systems can

also take inspiration from *Modular Robotics* to use small robotic blocks that can be snapped together and removed from the table [51]. My colleague Philipp Schoessler explored this idea further by developing techniques for arranging modular blocks on the *inFORM* shape display for his thesis on *Shape Synthesis* [175].

*Limited resolution* of current shape displays affect what content can be represented. Due to sampling requirements we need to have a resolution twice that of the smallest feature to be able to render it. We observed this issue when rendering fingers (2 cm wide) on the *TRANSFORM* shape display with 2.54 cm spacing. Increasing the resolution and pin travel will allow for more complex content and more realistic representation of people and objects.

*Noise* When moving at full speed, our shape displays produce up to 85dB at a distance of 1 meter. This can be distracting to users involved in conversation or work, but can be improved through different actuation hardware. We are also investigating using the sound as an additional dimension for performances.

*Network latency* was not investigated in depth on our current telepresence prototypes, as these were not deployed as distributed setups. However, the effect of latency is a critical factor for effective remote manipulation, and further studies are required to investigate how it will affect operator performance.

*Haptic Rendering* is limited by our current actuators. With faster, stronger actuators, we can simulate how different materials feel, or present varying resistance as users push into objects to simulate internal structures. This would be particularly interesting for medical scenarios. Besides using more precise, faster actuators, stiffness changing materials on top of the pins can add an additional haptic layer to the rendered shapes [36, 46].

*Applications:* Beyond these technical improvements to the system, we would also like to further explore other application domains, such as remote systems for education, medical, and industrial automation. We envision that teachers could use physically rendered models to discuss history and mathematics, remotely support students during physical tasks, or slow down the recorded playback of their hands so that pupils can understand their movements more clearly. Doctors could use *Physical*

*Telepresence* for remote palpation of patients, or to discuss volumetric medical data sets between distributed teams of experts. In industrial automation and robotics, shape displays can allow for unique transportation and parallel handling of various materials, and could be combined with conventional robotic manipulators for added control over objects.



# Chapter 8

## Conclusion

This dissertation has introduced a set of spatial interaction techniques for shape displays that range from touching the interface surface, to interacting with tangible objects on top, and to engaging through body gestures in relation to it. These techniques were implemented on our custom-built shape display systems that integrate physical rendering, synchronized visual display, shape sensing, and spatial tracking of gestures and objects. On top of this hardware platform, applications for computer-aided design, urban planning, volumetric data exploration, and remote collaboration allow users to manipulate data at different scales and modalities.

Physical shape change evokes a strong emotional reaction and can startle and delight users. This capability requires thoughtful motion design. Based on our experiences of building applications and public installations, we present a set of animation principles for controlling the speed, quality of motion, scale, and direction of shape change, and to augment it with graphics and interactivity. Through our work, we discovered that shape displays are more than haptic interfaces; similar to kinetic sculptures, they relate to the users body in space even before touching the surface. This starts when a user moves around the shape to view it from different angles, but also leads to completely new interactions when the shape reacts to the users body movement. Shape displays are also powerful tools for animators by recording improvised physical motions of gestures, objects, and material interactions, and to combine these with pre-scripted and procedural animations. Our proposed animation

tools open a rich area for designing shape transformations, and could be extended to support “Wizard of Oz” experiments to test and prototype future computer interfaces, or to allow end users to add their own expressive character to shape-changing objects. By capturing and rendering shapes remotely, shape displays can also bridge physical distances for remote collaboration through shared telepresence workspaces. Users can manipulate shared models and handle remote objects, while augmenting their capabilities through altered body representations.

The insights gained from this dissertation point towards a future of rich physical interaction enabled by shape changing materials. As our everyday interactions with computers move away from desktop computer screens and smartphones and into the physical environment, such materials will allow us to experience information with our body, to improvise with gestures and objects when interacting with computers, and to personalize how the objects around us transform and respond to us.

# Bibliography

- [1] Jason Alexander, Andrés Lucero, and Sriram Subramanian. Tilt displays: Designing display surfaces with multi-axis tilting and actuation. In *Proceedings of the 14th International Conference on Human-computer Interaction with Mobile Devices and Services*, MobileHCI '12, pages 161–170, New York, NY, USA, 2012. ACM.
- [2] ALPS. Motor-driven master type (motor n fader, motor k fader, motor v fader). [http://www.alps.com/prod/info/E/HTML/Potentiometer/SlidePotentiometers/RSN1M/RSN1M\\_list.html](http://www.alps.com/prod/info/E/HTML/Potentiometer/SlidePotentiometers/RSN1M/RSN1M_list.html), 2015.
- [3] Mahoro Anabuki and Hiroshi Ishii. Ar-jig: A handheld tangible user interface for modification of 3d digital form via 2d physical curve. In *ISMAR '07*, pages 1–10. IEEE Computer Society, 2007.
- [4] S. Audet and M. Okutomi. A user-friendly method to geometrically calibrate projector-camera systems. In *IEEE CVPR 2009*, pages 47–54, 2009.
- [5] C. Basdogan, C.-H. Ho, and M.A. Srinivasan. Virtual environments for medical training: graphical and haptic simulation of laparoscopic common bile duct exploration. *Mechatronics, IEEE/ASME Transactions on*, 6(3):269–285, Sep 2001.
- [6] Olivier Bau, Ivan Poupyrev, Ali Israr, and Chris Harrison. Teslatouch: Electro-vibration for touch surfaces. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology*, UIST '10, pages 283–292, New York, NY, USA, 2010. ACM.
- [7] Patrick Baudisch, Nathaniel Good, Victoria Bellotti, and Pamela Schraedley. Keeping things in context: A comparative evaluation of focus plus context screens, overviews, and zooming. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '02, pages 259–266, New York, NY, USA, 2002. ACM.
- [8] Mohamed Benali-khoudja, Moustapha Hafez, Jean marc Alex, and Abderrahmane Kheddar. Tactile interfaces: a state-of-the-art survey. In *Int. Symposium on Robotics*, pages 721–726, 2004.

- [9] Hrvoje Benko, Ricardo Jota, and Andrew Wilson. MirageTable: freehand interaction on a projected augmented reality tabletop. In *ACM CHI '12*, pages 199–208, 2012.
- [10] O. Bimber, B. Frohlich, D. Schmalstieg, and L.M. Encarnacao. The virtual showcase. *Computer Graphics and Applications, IEEE*, 21(6):48–55, nov/dec 2001.
- [11] Oliver Bimber and Ramesh Raskar. *Spatial augmented reality : merging real and virtual worlds*. A K Peters, Wellesley, Mass, 2005.
- [12] Matthew Blackshaw, Anthony DeVincenzi, David Lakatos, Daniel Leithinger, and Hiroshi Ishii. Recompose: Direct and gestural interaction with an actuated surface. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems, CHI EA '11*, pages 1237–1242, New York, NY, USA, 2011. ACM.
- [13] J.C. Bliss, M.H. Katcher, C.H. Rogers, and R.P. Shepard. Optical-to-tactile image conversion for the blind. *Man-Machine Systems, IEEE Transactions on*, 11(1):58–65, March 1970.
- [14] David Bowen. Tele-present wind. In *ACM SIGGRAPH 2011 Art Gallery, SIGGRAPH '11*, pages 358–359, New York, NY, USA, 2011. ACM.
- [15] Scott Brave, Hiroshi Ishii, and Andrew Dahley. Tangible interfaces for remote collaboration and communication. In *ACM CSCW '98*, pages 169–178, New York, New York, USA, November 1998.
- [16] Frederick P. Brooks, Jr., Ming Ouh-Young, James J. Batter, and P. Jerome Kilpatrick. Project gropehaptic displays for scientific visualization. *SIGGRAPH Comput. Graph.*, 24(4):177–185, September 1990.
- [17] Steve Bryson and Creon Levit. The virtual windtunnel: An environment for the exploration of three-dimensional unsteady flows. In *Proceedings of the 2Nd Conference on Visualization '91, VIS '91*, pages 17–24, Los Alamitos, CA, USA, 1991. IEEE Computer Society Press.
- [18] Vannevar Bush and Jingtao Wang. As we may think. *Atlantic Monthly*, 176:101–108, 1945.
- [19] Jean Cassou. *2 Kinetic Sculptors: Nicolas Schöffer and Jean Tinguely*. Jewish Museum, 1965.
- [20] Eames Charles and Eames Ray. Powers of ten: A film dealing with the relative size of things in the universe and the effect of adding another zero. IBM, 1977.
- [21] Craig Michae Ciesla and Micah B. Yairi. User interface system and method, 07 2009. US Patent 8,179,375.

- [22] Marcelo Coelho, Hiroshi Ishii, and Pattie Maes. Surfex: A programmable surface for the design of tangible interfaces. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '08, pages 3429–3434, New York, NY, USA, 2008. ACM.
- [23] Marcelo Coelho and Jamie Zigelbaum. Shape-changing interfaces. *Personal Ubiquitous Comput.*, 15(2):161–173, February 2011.
- [24] MarkR. Cutkosky and RobertD. Howe. Human grasp choice and robotic grasp analysis. In SubramanianT. Venkataraman and Thea Iberall, editors, *Dextrous Robot Hands*, pages 5–31. Springer New York, 1990.
- [25] Clark D. Della Silva. *TactiForm: A Touch Sensitive Actuated Shape Display as a Dynamic Interactive User Interface for Geospatial Visual Analytic Systems*. PhD thesis, Massachusetts Institute of Technology, 2015.
- [26] Tom Djajadiningrat, Ben Matthews, and Marcelle Stienstra. Easy doesn't do it: Skill and expression in tangible aesthetics. *Personal Ubiquitous Comput.*, 11(8):657–676, December 2007.
- [27] P. Dourish. *Where the Action is: The Foundations of Embodied Interaction*. Bradford books. Bradford Books, 2004.
- [28] Anthony Dunne and Fiona Raby. *Speculative Everything: Design, Fiction, and Social Dreaming*. The MIT Press, 2013.
- [29] Ebayink. Tablet use 1. <https://flic.kr/p/bomL7q>, 2011.
- [30] Katherine M. Everitt, Scott R. Klemmer, Robert Lee, and James A. Landay. Two worlds apart: bridging the gap between physical and virtual media for distributed design collaboration. In *ACM CHI '03*, pages 553–560, April 2003.
- [31] Gregg E. Favalora. Volumetric 3d displays and application infrastructure. *Computer*, 38(8):37–44, August 2005.
- [32] Festo. WaveHandling:Transporting and sorting in one. In [http://www.festo.com/cms/en\\_corp/13136.htm](http://www.festo.com/cms/en_corp/13136.htm), 2013.
- [33] Kenneth P. Fishkin. A taxonomy for and analysis of tangible interfaces. *Personal Ubiquitous Comput.*, 8(5):347–358, September 2004.
- [34] Ward Fleming. Pin screen, August 27 1985. US Patent 4,536,980.
- [35] James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics: Principles and Practice (2Nd Ed.)*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1990.
- [36] Sean Follmer, Daniel Leithinger, Alex Olwal, Nadia Cheng, and Hiroshi Ishii. Jamming user interfaces: programmable particle stiffness and sensing for malleable and shape-changing devices. In *UIST '12*, pages 519–528. ACM, 2012.

- [37] Sean Follmer, Daniel Leithinger, Alex Olwal, Akimitsu Hogge, and Hiroshi Ishii. inform: Dynamic physical affordances and constraints through shape and object actuation. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*, UIST '13, pages 417–426, New York, NY, USA, 2013. ACM.
- [38] Sean Weston Follmer. *Dynamic Physical Affordances for Shape-Changing and Deformable User Interfaces*. PhD thesis, Massachusetts Institute of Technology, 2015.
- [39] Joe Francica. Interview with xenotran founder, dr. derrick page. <http://www.directionsmag.com/entry/interview-with-xenotran-founder-dr-derrick-page/123593>, November 2004. Accessed: 2015-06-11.
- [40] Michael Fritschi, Marc O Ernst, and Martin Buss. Integration of kinesthetic and tactile display—a modular design concept. In *Proceedings of the EuroHaptics*, pages 607–612, 2006.
- [41] G. W. Furnas. Generalized fisheye views. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '86, pages 16–23, New York, NY, USA, 1986. ACM.
- [42] N. Gabo and A. Pevsner. *Das Realistische Manifest*. Berlinische Galerie, 1920.
- [43] N. Garcia-Hernandez, N.G. Tsagarakis, and D.G. Caldwell. Feeling through tactile displays: A study on the effect of the array density and size on the discrimination of tactile patterns. *Haptics, IEEE Transactions on*, 4(2):100–110, March 2011.
- [44] Joshua Gargus, Byungmoon Kim, Jaroslaw R Rossignac, and Christopher D Shaw. Finger sculpting with digital clay. Technical Report GIT-GVU-02-22, Georgia Institute of Technology, 2002.
- [45] William W. Gaver. Technology affordances. In *CHI '91*, pages 79–84. ACM, 1991.
- [46] Adam M. Genecov, Andrew A. Stanley, and Allison M. Okamura. Perception of a Haptic Jamming display: Just noticeable differences in stiffness and geometry. In *2014 IEEE Haptics Symposium (HAPTICS)*, pages 333–338. IEEE, February 2014.
- [47] James George. ofxtimeline, 2015. [Online; accessed 26-August-2015].
- [48] James J Gibson. Observations on active touch. *Psychological review*, 69(6):477, 1962.
- [49] James Jerome Gibson. *The Senses Considered as Perceptual Systems*. Houghton Mifflin Co., 1966.

- [50] James Jerome Gibson. *The ecological approach to visual perception*. Psychology Press, 1986.
- [51] Kyle Gilpin, Keith Kotay, and Daniela Rus. Miche: Modular Shape Formation by Self-Dissassembly. In *IEEE International Conference on Robotics and Automation 2007*, pages 2241–2247. IEEE, April 2007.
- [52] S.C. Goldstein, J.D. Campbell, and T.C. Mowry. Programmable matter. *Computer*, 38(6):99–101, May 2005.
- [53] Mark Burry Goulthorpe, M. and G. Dunlop. Aegis hyposurface: the bordering of university and practice. In *Proceedings of the Twenty First Annual Conference of the Association for Computer-Aided Design in Architecture*, pages 344–349. ACADIA, 2001.
- [54] Markus Gross, Silke Lang, Kai Strehlke, Andrew Vande Moere, Oliver Staadt, Stephan Würmlin, Martin Naef, Edouard Lamboray, Christian Spagno, Andreas Kunz, Esther Koller-Meier, Tomas Svoboda, and Luc Van Gool. blue-c: a spatially immersive display and 3D video portal for telepresence. *ACM Transactions on Graphics*, 22(3):819–827, July 2003.
- [55] Tovi Grossman and Ravin Balakrishnan. Collaborative interaction with volumetric displays. In *ACM CHI '08*, pages 383–392, April 2008.
- [56] Tovi Grossman and Ravin Balakrishnan. Collaborative interaction with volumetric displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '08*, pages 383–392, New York, NY, USA, 2008. ACM.
- [57] 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, UIST '04*, pages 61–70, New York, NY, USA, 2004. ACM.
- [58] Antal Haans and Wijnand IJsselsteijn. Mediated social touch: a review of current research and future directions. *Virtual Reality*, 9(2-3):149–159, 2006.
- [59] Martin Hachet, Benoit Bossavit, Aurélie Cohé, and Jean-Baptiste de la Rivière. Toucheo: multitouch and stereo combined in a seamless workspace. In *UIST '11*, pages 587–592, 2011.
- [60] John Hardy, Christian Weichel, Faisal Taher, John Vidler, and Jason Alexander. Shapeclip: Towards rapid prototyping with shape-changing displays for designers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15*, pages 19–28, New York, NY, USA, 2015. ACM.
- [61] Chris Harrison, Hrvoje Benko, and Andrew D. Wilson. Omnitouch: wearable multitouch interaction everywhere. In *UIST '11*, pages 441–450, 2011.

- [62] Chris Harrison and Scott E. Hudson. Providing dynamically changeable physical buttons on a visual display. In *CHI '09*, pages 299–308. ACM, 2009.
- [63] Rex Hartson. Cognitive, physical, sensory, and functional affordances in interaction design. *Behaviour & Information Technology*, 22(5):315–338, 2003.
- [64] Sunao Hashimoto, Akihiko Ishida, Masahiko Inami, and Takeo Igarashi. Touchme: An augmented reality based remote robot manipulation. In *The 21st International Conference on Artificial Reality and Telexistence, Proceedings of ICAT2011*, ICAT '11, 2011.
- [65] Jeffrey Heer and Ben Shneiderman. Interactive dynamics for visual analysis. *Commun. ACM*, 55(4):45–54, April 2012.
- [66] Steven J. Henderson and Steven Feiner. Opportunistic controls: leveraging natural affordances as tangible user interfaces for augmented reality. In *VRST '08*, pages 211–218, 2008.
- [67] Otmar Hilliges, David Kim, Shahram Izadi, Malte Weiss, and Andrew Wilson. Holodesk: direct 3d interactions with a situated see-through display. In *CHI '12*, pages 2421–2430, 2012.
- [68] W.D. Hillis and B. Ferren. Raised display apparatus, October 14 2008. US Patent 7,436,388.
- [69] Michael A. Hiltzik and Rutkoff. *Dealers of Lightning: Xerox PARC and the Dawn of the Computer Age*. HarperCollins Publishers, 1999.
- [70] K. Hirota and M. Hirose. Simulation and presentation of curved surface in virtual reality environment through surface display. In *Virtual Reality Annual International Symposium, 1995. Proceedings.*, pages 211–216, Mar 1995.
- [71] Daniel Hirschmann. Glowbits. In *Ubicomp 2004, Adjunct Proceedings*. ACM, 2004.
- [72] Eve Hoggan, Stephen A. Brewster, and Jody Johnston. Investigating the effectiveness of tactile feedback for mobile touchscreens. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pages 1573–1582, New York, NY, USA, 2008. ACM.
- [73] John M. Hollerbach, Ian W. Hunter, and John Ballantyne. The robotics review 2. chapter A Comparative Analysis of Actuator Technologies for Robotics, pages 299–342. MIT Press, Cambridge, MA, USA, 1992.
- [74] Oren Horev. *Talking to the Hand*. PhD thesis, Interaction Design Institute Ivrea, 2006.
- [75] Takayuki Hoshi, Masafumi Takahashi, Kei Nakatsuma, and Hiroyuki Shinoda. Touchable holography. In *SIGGRAPH '09*, pages 23:1–23:1. ACM, 2009.

- [76] Jason Huggins. Photograph of “pinscreen toy” by ward fleming. <https://flic.kr/p/atLQAB>, 2011.
- [77] Immersion. Cybertouch. <http://www.vrlogic.com/html/immersion/cybertouch.html>, 2015.
- [78] H Ishii, C Ratti, B Piper, Y Wang, A Biderman, and E Ben-Joseph. Bringing Clay and Sand into Digital Design — Continuous Tangible user Interfaces. *BT Technology Journal*, 22(4):287–299, October 2004.
- [79] Hiroshi Ishii and Minoru Kobayashi. ClearBoard: a seamless medium for shared drawing and conversation with eye contact. In *ACM CHI '92*, pages 525–532, June 1992.
- [80] Hiroshi Ishii, Dávid Lakatos, Leonardo Bonanni, and Jean-Baptiste Labrune. Radical atoms: beyond tangible bits, toward transformable materials. *interactions*, 19(1):38–51, January 2012.
- [81] Hiroshi Ishii, Daniel Leithinger, Sean Follmer, Amit Zoran, Philipp Schoessler, and Jared Counts. Transform: Embodiment of "radical atoms" at milano design week. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems, CHI EA '15*, pages 687–694, New York, NY, USA, 2015. ACM.
- [82] Hiroshi Ishii and Brygg Ullmer. Tangible bits: towards seamless interfaces between people, bits and atoms. In *CHI '97*, pages 234–241. ACM, 1997.
- [83] Kentaro Ishii, Shengdong Zhao, Masahiko Inami, Takeo Igarashi, and Michita Imai. Designing laser gesture interface for robot control. In *INTERACT '09*, pages 479–492, 2009.
- [84] Takayuki Iwamoto, Mari Tatezono, and Hiroyuki Shinoda. Non-contact method for producing tactile sensation using airborne ultrasound. In *Proceedings of the 6th International Conference on Haptics: Perception, Devices and Scenarios, EuroHaptics '08*, pages 504–513, Berlin, Heidelberg, 2008. Springer-Verlag.
- [85] Hiroo Iwata. Art and technology in interface devices. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology, VRST '05*, pages 1–7, New York, NY, USA, 2005. ACM.
- [86] Hiroo Iwata, Hiroaki Yano, Fumitaka Nakaizumi, and Ryo Kawamura. Project feelix: adding haptic surface to graphics. In *SIGGRAPH '01*, pages 469–476. ACM, 2001.
- [87] Hiroo Iwata, Hiroaki Yano, and Naoto Ono. Volflex. In *ACM SIGGRAPH 2005 Emerging Technologies, SIGGRAPH '05*, New York, NY, USA, 2005. ACM.

- [88] Robert J.K. Jacob, Audrey Girouard, Leanne M. Hirshfield, Michael S. Horn, Orit Shaer, Erin Treacy Solovey, and Jamie Zigelbaum. Reality-based interaction: A framework for post-wimp interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pages 201–210, New York, NY, USA, 2008. ACM.
- [89] Yvonne Jansen, Pierre Dragicevic, Petra Isenberg, Jason Alexander, Abhijit Karnik, Johan Kildal, Sriram Subramanian, and Kasper Hornbæk. Opportunities and challenges for data physicalization. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 3227–3236, New York, NY, USA, 2015. ACM.
- [90] N.Y.) Jewish Museum (New York, Smithsonian Institution, and American Motors Corporation. *Software: Information Technology : Its New Meaning for Art*. Jewish Museum, 1970.
- [91] Bradley Carter Kaanta. *PINS : a haptic computer interface system*. PhD thesis, Massachusetts Institute of Technology, 2004.
- [92] Kurt A. Kaczmarek and Paul Bach-Y-Rita. Virtual environments and advanced interface design. chapter Tactile Displays, pages 349–414. Oxford University Press, Inc., New York, NY, USA, 1995.
- [93] Hiroyuki Kajimoto, Yonezo Kanno, and Susumu Tachi. Forehead retina system. In *ACM SIGGRAPH 2006 Emerging Technologies*, SIGGRAPH '06, New York, NY, USA, 2006. ACM.
- [94] Victor Kaptelinin and Bonnie Nardi. Affordances in hci: toward a mediated action perspective. In *CHI '12*, pages 967–976. ACM, 2012.
- [95] Chang-Min Kim and Tek-Jin Nam. G-raff: An elevating tangible block for spatial tabletop interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 4161–4164, New York, NY, USA, 2015. ACM.
- [96] Seoktae Kim, Hyunjung Kim, Boram Lee, Tek-Jin Nam, and Woohun Lee. Inflatable mouse: Volume-adjustable mouse with air-pressure-sensitive input and haptic feedback. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pages 211–224, New York, NY, USA, 2008. ACM.
- [97] David Kirby. The future is now: Diegetic prototypes and the role of popular films in generating real-world technological development. *Social Studies of Science*, 2009.
- [98] Roberta L Klatzky, SJ Lederman, J Pellegrino, S Doherty, and B McClosky. Procedures for haptic object exploration vs. manipulation. *Vision and action: The control of grasping*, pages 110–127, 1990.

- [99] Scott R. Klemmer, Björn Hartmann, and Leila Takayama. How bodies matter: five themes for interaction design. In *DIS '06*, pages 140–149. ACM, 2006.
- [100] Sachiko Kodama. Dynamic ferrofluid sculpture: Organic shape-changing art forms. *Commun. ACM*, 51(6):79–81, June 2008.
- [101] Kurt Koffka. *Principles of Gestalt Psychology*. Harcourt Brace, New York, 1935.
- [102] Yohei Kojima, Taku Fujimoto, Yuichi Itoh, and Kosuke Nakajima. Polka dot: The garden of water spirits. In *SIGGRAPH Asia 2013 Emerging Technologies*, SA '13, pages 15:1–15:2, New York, NY, USA, 2013. ACM.
- [103] Lawrence M. Krauss. *The Physics of Star Trek*. Perseus Books, Cambridge, MA, USA, 2007.
- [104] T.W. Kuhlen and B. Hentschel. Quo vadis cave: Does immersive visualization still matter? *Computer Graphics and Applications, IEEE*, 34(5):14–21, Sept 2014.
- [105] John Lasseter. Principles of traditional animation applied to 3d computer animation. *SIGGRAPH Comput. Graph.*, 21(4):35–44, August 1987.
- [106] Alberto Laurenzi. Photograph of “strutturazione pulsante” (1983) by gianni colombo at galleria d’arte contemporanea di suzzara. <http://www.premiosuzzara.it/ura/Gianni-Colombo.html>, 2000.
- [107] Min Kyung Lee and Leila Takayama. "Now, i have a body": uses and social norms for mobile remote presence in the workplace. In *ACM CHI '11*, pages 33–42, May 2011.
- [108] Narae Lee, Ju-Whan Kim, Jungsoo Lee, Myeongsoo Shin, and Woohun Lee. MoleBot: a robotic creature based on physical transformability. In *VRIC '12*, pages 17:1–17:4, March 2012.
- [109] Sangyoon Lee, G. Sukhatme, Jounghyun Kim, and Chan-Mo Park. Haptic control of a mobile robot: a user study. In *Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on*, volume 3, pages 2867–2874 vol.3, 2002.
- [110] Vincent Lee, Amy Catherine Nau, Charles Laymon, Kevin Chan, Beda Rosario, and Chris Fisher. Successful tactile based visual sensory substitution use functions independently of visual pathway integrity. *Frontiers in Human Neuroscience*, 8(291), 2014.
- [111] Daniel Leithinger, Sean Follmer, Alex Olwal, and Hiroshi Ishii. Physical telepresence: shape capture and display for embodied, computer-mediated remote collaboration. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, UIST '14, pages 461–470, New York, NY, USA, 2014. ACM.

- [112] Daniel Leithinger, Sean Follmer, Alex Olwal, Samuel Luescher, Akimitsu Hogge, Jinha Lee, and Hiroshi Ishii. Sublimate: State-changing virtual and physical rendering to augment interaction with shape displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 1441–1450, New York, NY, USA, 2013. ACM.
- [113] Daniel Leithinger and Hiroshi Ishii. Relief: A scalable actuated shape display. In *Proceedings of the Fourth International Conference on Tangible, Embedded, and Embodied Interaction*, TEI '10, pages 221–222, New York, NY, USA, 2010. ACM.
- [114] Daniel Leithinger, David Lakatos, Anthony DeVincenzi, Matthew Blackshaw, and Hiroshi Ishii. Direct and gestural interaction with relief: a 2.5d shape display. In *UIST '11*, pages 541–548. ACM, 2011.
- [115] Mitchell JH Lum, Jacob Rosen, Thomas S Lendvay, Mika N Sinanan, and Blake Hannaford. Effect of time delay on telesurgical performance. In *IEEE ICRA '09*, pages 4246–4252, 2009.
- [116] Steve Mann. Telepointer: Hands-Free Completely Self Contained Wearable Visual Augmented Reality without Headwear and without any Infrastructural Reliance. In *2012 16th International Symposium on Wearable Computers*, pages 177–178, October 2000.
- [117] Nicolai Marquardt, Miguel A. Nacenta, James E. Young, Sheelagh Carpendale, Saul Greenberg, and Ehud Sharlin. The haptic tabletop puck: Tactile feedback for interactive tabletops. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, ITS '09, pages 85–92, New York, NY, USA, 2009. ACM.
- [118] Thomas H. Massie and Kenneth J. Salisbury. PHANToM haptic interface: a device for probing virtual objects. In *Proceedings of the 1994 International Mechanical Engineering Congress and Exposition*, volume 55-1 of *IMECE '94*, pages 295–299. ASME, 1994.
- [119] Malcolm McCullough. *Abstracting Craft: The Practiced Digital Hand*. MIT Press, Cambridge, MA, USA, 1996.
- [120] W. A. McNeely. Robotic graphics: A new approach to force feedback for virtual reality. In *Proceedings of the 1993 IEEE Virtual Reality Annual International Symposium*, VRAIS '93, pages 336–341, Washington, DC, USA, 1993. IEEE Computer Society.
- [121] G. Michelitsch, J. Williams, M. Osen, B. Jimenez, and S. Rapp. Haptic chameleon: a new concept of shape-changing user interface controls with force feedback. In *CHI EA '04*, pages 1305–1308. ACM, 2004.

- [122] Paul Milgram and Fumio Kishino. A taxonomy of mixed reality visual displays. *IEICE TRANSACTIONS on Information and Systems*, 77(12):1321–1329, 1994.
- [123] T. Mitsuda, S. Kuge, M. Wakabayashi, and S. Kawamura. Wearable haptic display by the use of a Particle Mechanical Constraint. *Proceedings 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. HAPTICS 2002*, pages 153–158, 2002.
- [124] J. Moloney. *Designing Kinetics for Architectural Facades: State Change*. Taylor & Francis, 2011.
- [125] Florian 'Floyd' Mueller, Luke Cole, Shannon O'Brien, and Wouter Walmlink. Airhockey over a distance. In *CHI '06 extended abstracts*, pages 1133–1138, April 2006.
- [126] Chris Munro and Daniel Walczyk. Reconfigurable Pin-Type Tooling: A Survey of Prior Art and Reduction to Practice. *Journal of Manufacturing Science and Engineering*, 129(3):551–565, January 2007.
- [127] M. Nakatani, H. Kajimoto, K. Vlack, D. Sekiguchi, N. Kawakami, and S. Tachi. Control method for a 3d form display with coil-type shape memory alloy. In *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pages 1332–1337, April 2005.
- [128] Masashi Nakatani, Hiroyuki Kajimoto, Naoki Kawakami, and Susumu Tachi. Tactile sensation with high-density pin-matrix. In *Proceedings of the 2Nd Symposium on Applied Perception in Graphics and Visualization, APGV '05*, pages 169–169, New York, NY, USA, 2005. ACM.
- [129] Masashi Nakatani, Hiroyuki Kajimoto, Dairoku Sekiguchi, Naoki Kawakami, and Susumu Tachi. 3d form display with shape memory alloy. In *ICAT*, 2003.
- [130] NASA. Ames developed (pop optics) goggles now at dulles annex of natl. air and space museum/smithsonian inst. [https://commons.wikimedia.org/wiki/File:AC89-0437-20\\_a.jpeg](https://commons.wikimedia.org/wiki/File:AC89-0437-20_a.jpeg), 2006.
- [131] William Newman. Must electronic gadgets disrupt our face-to-face conversations? *interactions*, 13(6):18–19, November 2006.
- [132] Ryuma Niiyama and Yoichiro Kawaguchi. Gemotion screen: A generative, emotional, interactive 3d display. In *Proc. ASIAGRAPH in Shanghai*, pages 115–120, Shanghai, China, July 2008.
- [133] Donald A. Norman. *The Psychology of Everyday Things*. Basic Books, Inc., 1988.

- [134] Ryo Oguchi, Yasuaki Kakehi, Keita Takahashi, and Takeshi Naemura. Photonastic surface: Pin matrix type display controlled with light. In *Proceedings of the 2008 International Conference on Advances in Computer Entertainment Technology*, ACE '08, pages 396–396, New York, NY, USA, 2008. ACM.
- [135] Alex Olwal, Christoffer Lindfors, Jonny Gustafsson, Torsten Kjellberg, and Lars Mattsson. Astor: An autostereoscopic optical see-through augmented reality system. In *ISMAR '05*, pages 24–27. IEEE Computer Society, 2005.
- [136] openFrameworks. openframeworks, 2015. [Online; accessed 26-August-2015].
- [137] Maria V. Ottermo, Åyvind Stavdahl, and Tor A. Johansen. A remote palpation instrument for laparoscopic surgery: Design and performance. *Minimally Invasive Therapy & Allied Technologies*, 18(5):259–272, 2009.
- [138] D.J. Page. Reconfigurable surface, June 7 2005. US Patent 6,903,871.
- [139] Gian Pangaro, Dan Maynes-Aminzade, and Hiroshi Ishii. The actuated workbench: computer-controlled actuation in tabletop tangible interfaces. In *UIST '02*, pages 181–190. ACM, 2002.
- [140] Amanda Parkes, Ivan Poupyrev, and Hiroshi Ishii. Designing kinetic interactions for organic user interfaces. *Communications of the ACM*, 51(6):58, June 2008.
- [141] James Patten and Hiroshi Ishii. Mechanical constraints as computational constraints in tabletop tangible interfaces. In *CHI '07*, pages 809–818. ACM, 2007.
- [142] Richard P Paul. *Robot manipulators: mathematics, programming, and control: the computer control of robot manipulators*. Richard Paul, 1981.
- [143] Eric Paulos and John Canny. PRoP: personal roving presence. In *ACM CHI '98*, pages 296–303, January 1998.
- [144] Esben Warming Pedersen and Kasper Hornbæk. Tangible bots: Interaction with active tangibles in tabletop interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 2975–2984, New York, NY, USA, 2011. ACM.
- [145] Ken Perlin and David Fox. Pad: An alternative approach to the computer interface. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '93, pages 57–64, New York, NY, USA, 1993. ACM.
- [146] Ben Piper and Hiroshi Ishii. Pegblocks: A learning aid for the elementary classroom. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '02, pages 686–687, New York, NY, USA, 2002. ACM.

- [147] Ben Piper, Carlo Ratti, and Hiroshi Ishii. Illuminating clay: A 3-d tangible interface for landscape analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '02, pages 355–362, New York, NY, USA, 2002. ACM.
- [148] Ben Piper, Carlo Ratti, and Hiroshi Ishii. Illuminating clay: a 3-d tangible interface for landscape analysis. In *CHI '02*, pages 355–362. ACM, 2002.
- [149] W.J. Plesniak, R.S. Pappu, and S.A. Benton. Haptic holography: a primitive computational plastic. *Proceedings of the IEEE*, 91(9):1443 – 1456, sept. 2003.
- [150] W.J. Plesniak, R.S. Pappu, and S.A. Benton. Haptic holography: a primitive computational plastic. *Proceedings of the IEEE*, 91(9):1443 – 1456, sept. 2003.
- [151] F. Popper. *Origins and development of kinetic art*. New York Graphic Society, 1968.
- [152] Ivan Poupyrev, Tatsushi Nashida, Shigeaki Maruyama, Jun Rekimoto, and Yasufumi Yamaji. Lumen: interactive visual and shape display for calm computing. In *ACM SIGGRAPH 2004 E Tech*, page 17, August 2004.
- [153] Ivan Poupyrev, Tatsushi Nashida, and Makoto Okabe. Actuation and Tangible User Interfaces: the Vaucanson Duck, Robots, and Shape Displays. In *ACM TEI '07*, pages 205–212, February 2007.
- [154] Denise Prescher, Gerhard Weber, and Martin Spindler. A tactile windowing system for blind users. In *Proceedings of the 12th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '10, pages 91–98, New York, NY, USA, 2010. ACM.
- [155] Hayes Solos Raffle, Amanda J. Parkes, and Hiroshi Ishii. Topobo: A constructive assembly system with kinetic memory. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '04, pages 647–654, New York, NY, USA, 2004. ACM.
- [156] Ramesh Raskar, Greg Welch, Matt Cutts, Adam Lake, Lev Stesin, and Henry Fuchs. The office of the future: a unified approach to image-based modeling and spatially immersive displays. In *ACM SIGGRAPH '98*, pages 179–188, July 1998.
- [157] Ramesh Raskar, Greg Welch, Kok-Lim Low, and Deepak Bandyopadhyay. Shader lamps: Animating real objects with image-based illumination. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, pages 89–102, London, UK, UK, 2001. Springer-Verlag.
- [158] Majken K. Rasmussen, Esben W. Pedersen, Marianne G. Petersen, and Kasper Hornbaek. Shape-changing interfaces: a review of the design space and open research questions. In *CHI '12*, pages 735–744. ACM, 2012.

- [159] Carlo Ratti, Yao Wang, Ben Piper, Hiroshi Ishii, and Assaf Biderman. Phoxel-space: an interface for exploring volumetric data with physical voxels. In *DIS '04*, pages 289–296, 2004.
- [160] Tim Regan, David Sweeney, John Helmes, Vasillis Vlachokyriakos, Siân Lindley, and Alex Taylor. Designing engaging data in communities. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems, CHI EA '15*, pages 271–274, New York, NY, USA, 2015. ACM.
- [161] Jan Richter, Bruce H. Thomas, Maki Sugimoto, and Masahiko Inami. Remote active tangible interactions. In *ACM TEI '07*, pages 39–42, February 2007.
- [162] George W. Rickey. The morphology of movement. *Art Journal*, 22(4):220–231, 1963.
- [163] Eckard Riedenklau, Thomas Hermann, and Helge Ritter. An integrated multi-modal actuated tangible user interface for distributed collaborative planning. In *ACM TEI '12*, pages 169–174, February 2012.
- [164] Daniel E Rivera, Manfred Morari, and Sigurd Skogestad. Internal model control: Pid controller design. *Industrial & engineering chemistry process design and development*, 25(1):252–265, 1986.
- [165] Stuart Pill Robert Young, Danial Pezzutti and Richard Sharp. The language of motion in industrial design. In *Design and semantics of form and movement*, pages 6–12, November 2005.
- [166] Diana Robinson. Child’s hands in clay. <https://flic.kr/p/jnpKNk>, 2013.
- [167] John W. Romanishin, Kyle Gilpin, and Daniela Rus. M-blocks: Momentum-driven, magnetic modular robots. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4288–4295. IEEE, November 2013.
- [168] Jarek Rossignac, Mark Allen, Wayne J. Book, Ari Glezer, Imme Ebert-Uphoff, Chris Shaw, David Rosen, Stephen Askins, Jing Bai, Paul Bosscher, Joshua Gargus, ByungMoon Kim, Ignacio Llamas, Austina Nguyen, Guang Yuan, and Haihong Zhu. Finger sculpting with digital clay: 3d shape input and output through a computer-controlled real surface. In *Proceedings of the Shape Modeling International 2003, SMI '03*, pages 229–, Washington, DC, USA, 2003. IEEE Computer Society.
- [169] Anne Roudaut, Abhijit Karnik, Markus Löchtefeld, and Sriram Subramanian. Morphees: Toward high "shape resolution" in self-actuated flexible mobile devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '13*, pages 593–602, New York, NY, USA, 2013. ACM.
- [170] Daniel Rozin. Wooden mirror. *IEEE Spectr.*, 38(3):69–, March 2001.

- [171] Emanuele Ruffaldi. Immersive haptic workstation at the virtual reality lab at epfl in lousanna (switzerland) directed by prof. daniel thalman. <http://teslacore.blogspot.com/2005/06/epfl-haptic-workstation.html>, 2005.
- [172] Joachim Sauter, Susanne Jaschko, and Jussi Ängeslevä. *ART+COM: Media Spaces and Installations*. Gestalten, Berlin, 2011.
- [173] Chris Scharver, Ray Evenhouse, Andrew Johnson, and Jason Leigh. Designing cranial implants in a haptic augmented reality environment. *Commun. ACM*, 47(8):32–38, August 2004.
- [174] Christopher Schmandt. Spatial input/display correspondence in a stereoscopic computer graphic work station. In *SIGGRAPH '83*, pages 253–261. ACM, 1983.
- [175] Philipp Schoessler. *Shape Synthesis: Physical Object Augmentation and Actuation for Display and Interaction on Shape Changing Interfaces*. PhD thesis, Massachusetts Institute of Technology, 2015.
- [176] Orit Shaer and Eva Hornecker. Tangible user interfaces: Past, present, and future directions. *Found. Trends Hum.-Comput. Interact.*, 3(1&#8211;2):1–137, January 2010.
- [177] Orit Shaer, Nancy Leland, Eduardo H. Calvillo-Gamez, and Robert J. K. Jacob. The tac paradigm: Specifying tangible user interfaces. *Personal Ubiquitous Comput.*, 8(5):359–369, September 2004.
- [178] Willoughby Sharp. *Air Art*. Kineticism Press, New York, NY, USA, 1966.
- [179] Nathan Shedroff and Christopher Noessel. *Make It So: Interaction Design Lessons from Science Fiction*. Rosenfeld Media, 2012.
- [180] Taiki Shimba and Joo-Ho Lee. Shadow-free interaction: A proposal for rear projection based shape display. *Procedia Technology*, 18(0):140 – 144, 2014. International workshop on Innovations in Information and Communication Science and Technology, {IICST} 2014, 3-5 September 2014, Warsaw, Poland.
- [181] B. (Director) Singer. X-men. 20th Century Fox, 2000.
- [182] C. Sjostrom. Designing haptic computer interfaces for blind people. In *Signal Processing and its Applications, Sixth International, Symposium on. 2001*, volume 1, pages 68–71 vol.1, 2001.
- [183] Lisa Skedung, Martin Arvidsson, Jun Young Chung, Christopher M. Stafford, Birgitta Berglund, and Mark W. Rutland. Feeling small: Exploring the tactile perception limits. *Sci. Rep.*, 3, 09 2013.
- [184] M. Smith and J.J. Clarke. *Stelarc: The Monograph*. Electronic Culture Series. MIT Press, 2005.

- [185] Rajinder Sodhi, Ivan Poupyrev, Matthew Glisson, and Ali Israr. Areal: Interactive tactile experiences in free air. *ACM Trans. Graph.*, 32(4):134:1–134:10, July 2013.
- [186] Rajinder S. Sodhi, Brett R. Jones, David Forsyth, Brian P. Bailey, and Giuliano Maciucci. BeThere: 3D mobile collaboration with spatial input. In *ACM CHI '13*, pages 179–188, April 2013.
- [187] Martin Spindler, Christian Tominski, Heidrun Schumann, and Raimund Dachsel. Tangible views for information visualization. In *ACM International Conference on Interactive Tabletops and Surfaces, ITS '10*, pages 157–166, New York, NY, USA, 2010. ACM.
- [188] Christoph Staub, Keita Ono, Hermann Mayer, Alois Knoll, Heinz Ulbrich, and Robert Bauernschmitt. Remote minimally invasive surgery –haptic feedback and selective automation in medical robotics. *Applied Bionics and Biomechanics*, 8(2):221–236, April 2011.
- [189] Duncan R. Stevenson, Kevin A. Smith, John P. McLaughlin, Chris J. Gunn, J. P. Veldkamp, and Mark J. Dixon. Haptic workbench: a multisensory virtual environment. *Proc. SPIE*, 3639:356–366, 1999.
- [190] Liz Stinson. Watch 800 wooden balls shape-shift into a perfect spiral. <http://www.wired.com/2014/08/watch-804-wooden-balls-shape-shift-into-a-perfect-spiral>, August 2014. Accessed: 2014-08-04.
- [191] Ivan E. Sutherland. The ultimate display. In *Proceedings of the Congress of the International Federation of Information Processing (IFIP)*, volume volume 2, pages 506–508, 1965.
- [192] Ivan E. Sutherland. A head-mounted three dimensional display. In *Proceedings of the December 9-11, 1968, Fall Joint Computer Conference, Part I, AFIPS '68 (Fall, part I)*, pages 757–764, New York, NY, USA, 1968. ACM.
- [193] S. Tachi, T. Maeda, R. Hirata, and H. Hoshino. A Construction Method of Virtual Haptic Space. In *Proceedings of the 4th International Conference on Artificial Reality and Tele-Existence (ICAT'94)*, 1994.
- [194] Faisal Taher, John Hardy, Abhijit Karnik, Christian Weichel, Yvonne Jansen, Kasper Hornbæk, and Jason Alexander. Exploring interactions with physically dynamic bar charts. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI '15*, pages 3237–3246, New York, NY, USA, 2015. ACM.
- [195] Shohei Takei, Makoto Iida, and Takeshi Naemura. Kinereels: Extension actuators for dynamic 3d shape. In *ACM SIGGRAPH 2011 Posters*, SIGGRAPH '11, pages 84:1–84:1, New York, NY, USA, 2011. ACM.

- [196] Joshua Tanenbaum. Design fictional interactions: Why hci should care about stories. *interactions*, 21(5):22–23, September 2014.
- [197] John C. Tang and Scott Minneman. VideoWhiteboard: video shadows to support remote collaboration. In *ACM CHI '91*, pages 315–322, March 1991.
- [198] John C. Tang and Scott L. Minneman. Videodraw: a video interface for collaborative drawing. *ACM Transactions on Information Systems*, 9(2):170–184, April 1991.
- [199] Masayuki Tani, Kimiya Yamaashi, Koichiro Tanikoshi, Masayasu Futakawa, and Shinya Tanifuji. Object-oriented video: interaction with real-world objects through live video. In *CHI '92*, pages 593–598. ACM, 1992.
- [200] Tommaso Toffoli and Norman Margolus. Programmable matter: Concepts and realization. In *Proceedings of the NATO Advanced Research Workshop on Lattice Gas Methods for PDE's : Theory, Applications and Hardware: Theory, Applications and Hardware*, pages 263–272, Amsterdam, The Netherlands, The Netherlands, 1991. North-Holland Publishing Co.
- [201] Jonas Togler, Fabian Hemmert, and Reto Wettach. Living interfaces: The thrifty faucet. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*, TEI '09, pages 43–44, New York, NY, USA, 2009. ACM.
- [202] Bruce Tognazzini. The “starfire” video prototype project: A case history. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '94, pages 99–105, New York, NY, USA, 1994. ACM.
- [203] Jessica Tsimeris, Colin Dedman, Michael Broughton, and Tom Gedeon. Forceform: A dynamically deformable interactive surface. In *Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces*, ITS '13, pages 175–178, New York, NY, USA, 2013. ACM.
- [204] Jessica Tsimeris, Duncan Stevenson, Matt Adcock, Tom Gedeon, and Michael Broughton. User created tangible controls using forceform: A dynamically deformable interactive surface. In *Proceedings of the Adjunct Publication of the 26th Annual ACM Symposium on User Interface Software and Technology*, UIST '13 Adjunct, pages 95–96, New York, NY, USA, 2013. ACM.
- [205] Jessica Tsimeris, Duncan Stevenson, Tom Gedeon, and Matt Adcock. Using forceform, a dynamically deformable interactive surface, for palpation simulation in medical scenarios. In *Proceedings of the Second International Workshop on Smart Material Interfaces: Another Step to a Material Future*, SMI '13, pages 19–22, New York, NY, USA, 2013. ACM.
- [206] Katherine M. Tsui, Munjal Desai, Holly A. Yanco, and Chris Uhlik. Exploring use cases for telepresence robots. In *ACM/IEEE HRI '11*, pages 11–18, March 2011.

- [207] B. Ullmer and H. Ishii. Emerging frameworks for tangible user interfaces. *IBM Syst. J.*, 39(3-4):915–931, July 2000.
- [208] Brygg Ullmer and Hiroshi Ishii. The metadesk: Models and prototypes for tangible user interfaces. In *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology*, UIST '97, pages 223–232, New York, NY, USA, 1997. ACM.
- [209] Brygg Ullmer, Hiroshi Ishii, and Robert J. K. Jacob. Token+constraint systems for tangible interaction with digital information. *ACM Trans. Comput.-Hum. Interact.*, 12(1):81–118, March 2005.
- [210] Brygg Ullmer, Hiroshi Ishii, and Robert J. K. Jacob. Token+constraint systems for tangible interaction with digital information. *ACM Trans. Comput.-Hum. Interact.*, 12(1):81–118, March 2005.
- [211] John Underkoffler and Hiroshi Ishii. Urp: A luminous-tangible workbench for urban planning and design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '99, pages 386–393, New York, NY, USA, 1999. ACM.
- [212] Peter Uray, DI Thomas Kienzl, and DI Ulf Marsche. Mri: A mixed reality interface for the masses. In *ACM SIGGRAPH 2006 Emerging Technologies*, SIGGRAPH '06, New York, NY, USA, 2006. ACM.
- [213] Leslie Carlson Vaughan. Understanding movement. In *CHI '97*, pages 548–549. ACM, 1997.
- [214] Bret Victor. A brief rant on the future of interaction design. <http://worrydream.com/ABriefRantOnTheFutureOfInteractionDesign/>, Nov 2011.
- [215] F. Vidal-Verdu and M. Hafez. Graphical tactile displays for visually-impaired people. *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, 15(1):119–130, March 2007.
- [216] Luke Vink, Viirj Kan, Ken Nakagaki, Daniel Leithinger, Sean Follmer, Philipp Schoessler, Amit Zoran, and Hiroshi Ishii. Transform as adaptive and dynamic furniture. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA '15, pages 183–183, New York, NY, USA, 2015. ACM.
- [217] Thorsten Völkel, Gerhard Weber, and Ulrich Baumann. Tactile graphics revised: The novel brailledis 9000 pin-matrix device with multitouch input. In *Proceedings of the 11th International Conference on Computers Helping People with Special Needs*, ICCHP '08, pages 835–842, Berlin, Heidelberg, 2008. Springer-Verlag.

- [218] Christian Weichel, Jason Alexander, and John Hardy. Shape display shader language (sdsl): A new programming model for shape changing displays. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA '15, pages 1121–1126, New York, NY, USA, 2015. ACM.
- [219] M. Weiser, R. Gold, and J.S. Brown. The origins of ubiquitous computing research at parc in the late 1980s. *IBM Systems Journal*, 38(4):693–696, 1999.
- [220] Norbert Wiener. *Cybernetics or Control and Communication in the Animal and the Machine*, volume 25. MIT press, 1961.
- [221] Wikipedia. Autodesk 3ds max — wikipedia, the free encyclopedia, 2015. [Online; accessed 26-August-2015].
- [222] Andrew D. Wilson. Robust computer vision-based detection of pinching for one and two-handed gesture input. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology*, UIST '06, pages 255–258, New York, NY, USA, 2006. ACM.
- [223] Andrew D Wilson and Daniel C Robbins. Playtogether: Playing games across multiple interactive tabletops. In *IUI Workshop on Tangible Play: Research and Design for Tangible and Tabletop Games*, 2007.
- [224] Faye Y Wu and Harry Asada. Bio-artificial synergies for grasp posture control of supernumerary robotic fingers. In *Robotics: Science and Systems X*, 2014.
- [225] Svetlana Yarosh, Anthony Tang, Sanika Mokashi, and Gregory D. Abowd. "almost touching": parent-child remote communication using the sharetable system. In *ACM CSCW '13*, pages 181–192, February 2013.
- [226] M. Yim, Wei-Min Shen, B. Salemi, D. Rus, M. Moll, H. Lipson, E. Klavins, and G.S. Chirikjian. Modular self-reconfigurable robot systems [grand challenges of robotics]. *Robotics Automation Magazine, IEEE*, 14(1):43–52, March 2007.
- [227] Takumi Yoshida, Sho Kamuro, Kouta Minamizawa, Hideaki Nii, and Susumu Tachi. Repro3d: full-parallax 3d display using retro-reflective projection technology. In *SIGGRAPH '10*, pages 20:1–20:1. ACM, 2010.
- [228] Haihong Zhu and Wayne J Book. Practical structure design and control for digital clay. In *ASME 2004 International Mechanical Engineering Congress and Exposition*, pages 1051–1058. American Society of Mechanical Engineers, 2004.
- [229] Jamie Zigelbaum, Alan Browning, Daniel Leithinger, Olivier Bau, and Hiroshi Ishii. G-stalt: A chirocentric, spatiotemporal, and telekinetic gestural interface. In *Proceedings of the Fourth International Conference on Tangible, Embedded, and Embodied Interaction*, TEI '10, pages 261–264, New York, NY, USA, 2010. ACM.