Magic Pen: A Versatile Digital Manipulative for Learning

by

Soheil Kianzad

B.Sc., K.N.Toosi University of Technology, 2012

M.Sc., The University of British Columbia (Vancouver), 2015

A THESIS SUBMITTED IN PARTIAL FULFILLMENT

OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES

(Computer Science)

The University of British Columbia

(Vancouver)

December 2021

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the dissertation entitled:

MAGIC PEN: A VERSATILE DIGITAL MANIPULATIVE FOR LEARNING

submitted by Soheil Kianzad in partial fulfillment of the requirements for the degree of DOCTOR OF PHILOSOPHY
in COMPUTER SCIENCE

Examining Committee:

KARON MACLEAN, PROFESSOR, DEPARTMENT OF COMPUTER SCIENCE, UBC
Supervisor

IVAN BESCHASTNIKH, ASSOCIATE PROFESSOR, DEPARTMENT OF COMPUTER SCIENCE, UBC
Supervisory Committee Member

MICHIEL VAN DE PANNE, PROFESSOR, DEPARTMENT OF COMPUTER SCIENCE, UBC
Supervisory Committee Member

JIM LITTLE, PROFESSOR, DEPARTMENT OF COMPUTER SCIENCE, UBC
University Examiner

MACHIEL VAN DER LOOS, ASSOCIATE PROFESSOR, DEPARTMENT OF MECHANICAL ENGINEERING, UBC
University Examiner

Additional Supervisory Committee Members:

PIERRE DILLENBOURG, PROFESSOR, SCHOOL OF COMPUTER & COMMUNICATION SCIENCES, EPFL

Supervisory Committee Member

Abstract

Digital manipulatives such as robots are an opportunity for interactive and engaging learning activities. The addition of haptic and specifically force feedback to digital manipulatives can enrich the learning of science-related concepts by building physical intuition. As a result, learners can design experiments and physically explore them to solve problems they have posed.

In my thesis, I present the evolution of the design and evaluation of a versatile digital manipulative – called MagicPen – in a human-centered design context. First, I investigate how force feedback can enable learners to fluidly express their ideas. I identify three core interactions as bases for physically assisted sketching (phasking). Then, I show how using these interactions improves the accuracy of users' drawings as well as their authority in creative works. In the next phase, I demonstrate the potential benefits of using force feedback in a collaborative learning framework, in a manner that is generalizable beyond the device we invented and lends insight on how haptics can empower digital manipulatives to express advanced concept by means of the behaviour of a virtual avatar and the respective feeling of force feedback. This informs our device's capability for learning advanced concepts in classroom settings and further considerations for the next iterations of the MagicPen.

Based on the findings of how haptic feedback could assist with design and exploration in learning, In the last phase of my thesis, I propose a framework for physically assisted learning (PAL) which links the expression and exploration of an idea. Furthermore, I explain how to instantiate the PAL framework using available technologies and discuss a path forward to a larger vision of physically assisted learning. PAL highlights the role of haptics in future "objects-to-think-with".

Lay Summary

Educational haptic platforms exploit various modalities to create effective interactive environments that can support embodied physical interactions. These platforms have the potential to leverage a student's physical intuition to make abstract topics in physics, math, and other fields of science more concrete. This project aims to create a versatile educational robot that serves as an object-to-think-with. It can provide students with intuitive ways to experience various science, technology, engineering, and mathematics (STEM) concepts by making them more accessible or helping students approach these concepts in a more tangible way. We will explain the evolution of the design and evaluation of our proposed device in a human-centered design context.

Preface

All the research outputs are the result of collaborative efforts as no creative work belongs to an individual. Here, I will clarify my contribution to the published works.

Chapter 3: Device Design

> S. Kianzad and K. E. MacLean, "Harold's purple crayon rendered in haptics: Large-stroke, handheld ballpoint force feedback," 2018 IEEE Haptics Symposium (HAPTICS), San Francisco, CA, 2018, pp. 106-111.

My supervisor (Dr. Karon MacLean) came up with the idea of drawing and then feeling it, besides the initial approach of how to implement it. I built upon my supervisor's idea and came up with a different and more realizable approach for implementing the mechatronics system. I contributed all the engineering work, design iterations, and mechanical tests with feedback and guidance from my supervisor. I wrote the first draft of the paper which was fully edited and re-written by my supervisor before the submission.

> S. Kianzad and K. E. MacLean, "Ballpoint Drive Haptic Force-Feedback Display", provisional patent disclosure.

I made all the drawings to convey our invention with feedback and guidance from my supervisor. My supervisor and I collaboratively shaped the framework, and I added the technical details. She re-wrote and revised the final patent draft.

Chapter 4: PHysically Assisted SKetching

> Soheil Kianzad, Yuxiang Huang, Robert Xiao, and Karon E. MacLean. 2020. Phasking on Paper: Accessing a Continuum of PHysically Assisted SKetchING. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20). Association for Computing Machinery, New York, NY, USA.

I designed the mechatronics system, wrote the state-model, and designed a closed-loop controller. I supervised Yuxiang Huang, an undergrad volunteer who built a lightweight CAD platform for the RPi computer. Our group collaborator, Dr. Xiao, helped with implementing the hardware PWM to control motors and the integration of a digital pen to achieve robust absolute tracking. I designed the user study with the help of Dr. Xiao and conducted the user-tests. I designed the experimental setup and ran mechanical tests on the device. My supervisor framed the paper and led the writing.

Chapter 5: Benefits of Force Feedback for Grounding

> Soheil Kianzad, Julia A. B. Lindsay, Wafa Johal, Unma Mayur Desai, Hala Khodr, Hsin Yun(Tiffany) Wu, Pierre Dillenbourg, Karon E. MacLean "Dialectic Touch: Exploiting Force Feedback for Grounding in Collaborative Learning Tasks", IEEE Transaction on Haptics (in preparation).

I designed the experiments with three haptic/robotic platforms collaboratively with my supervisor and Julia Lindsay– an undergrad student that I was supervising. I programmed and implemented all these experiments. I ran pilot tests and iterated on the design for each study. Julia designed the pre-test and post-tests for the three studies. She also drafted an amendment for conducting our study during the pandemic. Together we designed and ran the user evaluations for the two studies. I conducted the quantitative analysis on our user study's result. I led the qualitative assessment while I received help from Unma Desai (master student) and Tiffany Wu (volunteer) to avoid potential biases. Unma, Tiffany and I individually performed the qualitative coding, and collaboratively, we discussed the findings. I received insightful feedback and comments from my supervisor, Dr. Wafa Johal, and Dr. Pierre Dillenbourg through out the process, specifically on the design of the experiments and running the analysis. I wrote the original draft and my supervisor did the final edits.

> S. Kianzad and K. E. MacLean, "Collaborating Through Magic Pens: Grounded Forces in Large, Overlappable Workspaces." International AsiaHaptics conference. Springer, Singapore, 2018.

A single-user configuration with Virtual Electrostatic Lab was initially developed by Lotus Hanzi Zhang in collaboration with Matthew Chun and myself for Student Innovation Challenge – IEEE World Haptics 2017. Later, I extended this framework to enable multi-user interaction and added more functionality to it. I made the haptic devices, developed the virtual Jigsaw puzzle, and drafted the demo paper [114] which was revised by my supervisor.

Chapter 6: The Physically Assisted Learning (PAL)

> Soheil Kianzad, Guanxiong Chan, Karon E. MacLean "PAL: A Framework for Physically Assisted Learning through Design and Exploration with a Haptic Robot Buddy," Frontiers in Robotics and AI, pp 228-250, Vol 8, 2021[117].

Based on my supervisor's initial vision and with her guidance, I proposed the PAL framework. I implemented the idea, and I came up with two examples: (1) handwriting and (2) mass-spring to demonstrate the concept. Further, I designed the passivity controller, and I show how to expand the idea to other domains of physics. I supervised Guanxiong Chan, an undergrad volunteer who helped with the literature review and Wifi communication. My supervisor had a major contribution to the writing of the paper.

All research including human participants were reviewed and approved by UBC's Behavioural Research Ethics Board, approval ID (H14-01763). This specifically includes the studies reported in Chapters 4 and 5.

*This project is a result of collaborative work between my supervisor and me, referred to as **"we"** throughout this document.*

Table of Contents

List of Tables

List of Figures

Acknowledgments

First and foremost, I would like to thank my Ph.D. supervisor, Dr. Karon E. MacLean, for all the supports and research mentorship that I received during these years. I highly appreciate her patience and encouragement for research excellence and challenging my work by urging better clarity in explaining my ideas. Her guidance was essential for shaping this work.

I am very grateful to my supervisory committee, Dr. Ivan Beschastnikh, Dr. Pierre Dillenbourg, and Dr. Michiel van de Panne, for their valuable feedback. I would like to thank my external examiner Dr. J. Edward Colgate, as well as my internal university examiners, Dr. Jim Little and Dr. Machiel Van der Loos for their insightful comments that led to several key improvements of this dissertation.

I thank Dr. Wafa Johal for assisting me to receive the Swiss National Centre of Competence in Research (NCCR)'s robotics fellowship and spending 4 months at the CHILI lab at EPFL. I learned many things from her.

I like to give special thanks to my student collaborators. In particular, Julia Lindsay, Yuxiang Huang, Guanxiong Chen, and Unma Desai for all our productive discussions and teamwork. My special thanks to my fellows at the SPIN lab (members and alums) for their supports and valuable input throughout this journey, Dr. Oliver Schnider, Dr. Hasti Seifi, Laura Cang, Paul Bucci, Matthew Chun, Salma Kashani, Dilan Ustek, Hanieh Shakeri, Preeti Vyas, Hannah Elbaggari, Rubia Guerra, Devyani McLaren, and Elizabeth Reid.

I thank the MUX group's faculty members and the students, for all their great feedback on my research and paper drafts. In particular, I thank the faculty members, Dr. Robert Xiao, Dr. Joanna McGrenere, Dr. Dongwook Yoon, Dr. Tamara Munzner, and Dr. Kellogg Booth, and the MUX students, Dr. Jessalyn Alvina, Dr. Francesco

Vitale. Yelim Kim, Taslim Arefin, Izabelle Janzen.

I would like to thank the CHILI lab members for their hospitality and my great time with them. In particular, Sina Shahmoradi, Dr. Thibault Lucien Christian Asselborn, Hala Khodr, Dr. Barbara Bruno, Dr. Arzu Güneysu Özgür, Utku Norman, and Jauwairia Nasir.

I want to thank the Natural Sciences and Engineering Council of Canada (NSERC) and UBC's Institute for Computing, Information and Cognitive Systems (ICICS), and the UBC Designing for People (DFP) Research Cluster for providing facilities and funding for this work. I thank the Swiss National Centre of Competence in Research Robotics for partially funding my visit to the CHILI Lab at EPFL, and in particular, Florence Colomb for organizing this visit. I thank all the participants who took part in our experiments.

Lastly, I have immense appreciation and love for my parents, my wife, my siblings, and my friends.

Dedication

This work is dedicated to my loving family, and to all the motivated and curious children with big dreams.

Chapter 1

Introduction

According to Papert [174], education should be more about exploring and engaging, with less emphasis on explaining. As discussed in the pedagogical theory of Constructionism, pupils should be involved in their process of learning, from designing and constructing meaningful projects/artifacts to exploring them and creating personal experiences. As one of the well-grounded theories in the domain of educational robotics, learners shape their knowledge based on what they already know, the experiences they gain and the way they organize this to construct knowledge [152]. In this chapter, we will propose an approach for making a useful technology to promote experiential learning based on the Constructionism learning theory, and outline how this dissertation will go about the design and validation of it. Specifically, we focus on empowering digital manipulatives with the addition of haptics and study their potential benefits.

Physical manipulatives are objects that aid learners in perceiving Physics and Math concepts by manipulating them. Pattern blocks, coloured chips, and coins are examples of physical manipulatives that are used in early childhood education to engage learners in hands-on activities. Digital manipulatives are physical objects with computational and communication capabilities that can promote different types of thinking in children by engaging them in playing and building. The history of using Digital manipulatives for education dates backs to the early 1970s - 1980s in several works from MIT, most notably from the Media Lab's Tangible Media and Epistemology and Learning groups and the Artificial Intelligence Lab. Among them, projects such as Floor Turtle [4], Graphical Logo, LEGO Mindstorms [174], Crickets, and Curlybot [64] introduced engaging environments to develop new approaches to thinking about mathematical concepts with encouraging results.

Robots are a class of digital manipulatives that use motion along with other visual or audio cues to express information. Children can program robots and therefore observe and experience how defining a set of rules results in predictable robot behaviours. This also gives them the freedom to decide what the robot is, based on how the robot behaves. However, as children grow older, the lack of versatility reduces the effectiveness of these robots in conveying advanced topics. Consequently, as teaching shifts more towards using formal methods in later years, they gradually disappear from classrooms.

We propose two exemplary difficulties of fitting this approach into conventional classroom settings, each pointing to a class of considerations.

- *School logistic considerations:* Teachers who are the major stakeholders of education have needs which are not met. For example, what would be a proper educative tool to help teachers deliver instructions and orchestrate the classroom more efficiently?

- *Tool versatility:* As more advanced topics are introduced to learners, software and hardware limitations make these robots less effective.We can name two main reasons behind this incompetency: (a) the tool does not permit learners to actively design, make, and change their learning environment based on their hypothesis (user's expressivity), and (b) it imposes a complicated level of interpretation to show the relation between the robot's motion and teaching subjects (tool's expressivity).

Even if these robots could not find their way to schools, that does not make them inferior digital manipulatives. In fact, there are clear pieces of evidence from school camps and workshops to show children learn several advanced topics just by playing with them [162]. One method that these robots use to improve the learning process is to help learners imagine themselves as the robot and further reason about the robot's motions based on the learners' commands. Accordingly, if these commands are given to solve a mathematical problem, then learners are provided with an intuitive way of understanding the Math concept by observing the robot's movements and reason about them [193]. Often, learners write code to communicate with robots; thus learners naturally get involved in systematic thinking and learn the key principles of computer science [229].

Despite the possible advantages of using robots to teach computer literacy, there is little conclusive evidence to show transferable learning from practicing computer programming with robots to learning other fields of science such as Physics and Math [73]. Moreover, there is a slight chance that using robot movements in some learning scenarios could be misleading, more possibly, when there is no direct link between the movements and information that needs to be delivered. Both cases optimistically expect students to either be able to apply the learned skills with robots to other scenarios (*e.g.*, using problem-solving skills such as divide-and-conquer in Math problems) or use their imagination to find the analogy between the robots' movements and Physics or Math concepts. In reality, this gap is often so large that learners cannot find the connection between the activity with robots and what they are supposed to learn.

We posit that combining the visual and haptic (sense of touch) cues can potentially address this deficit. By using the visual cues we obtain a proper environment that directly matches the learning concepts. With a haptically augmented manipulative, haptic cues build on the visual input available from a robot's motion as the user manipulates the virtual environment by moving and interacting with it through the handheld robot. This can provide complementary information about a represented concept, reduce cognitive load of unimodal input, and generally make the manipulative more expressive than one which works through vision alone as it is used to explore advanced STEM concepts. The combination of these two modalities reduces the risk of misinterpretation.

In this dissertation, we make three observations, the joint addressing of which provides a unique opportunity in this as well as some other fields.

1. Assisted sketching can improve fluid expression of ideas. Proper assistive force feedback to users' pens while sketching can help them to manifest and communicate their ideas to other people and to a computer.

2. Exploring the role of force feedback in learning in a human centered design

context could lead to understanding the best strategies for employing them in learning advanced topics in Physics and Math.

3. The addition of haptic feedback to a digital manipulative (beyond, for example, the ability to program motion alone) is potentially helpful to make more compelling interpretations so that learners can predict and reason about the outcome, based on what they see and what they feel. As a result, we can exploit this design space to empower learners to actively design, make, and change their learning environment based on the learners' hypothesis.

Countering these opportunities, we also observe the difficulties of incorporating haptic feedback into a digital manipulative suitable for a classroom environment. First, we could not find a mechanism to support low-cost untethered haptic feedback for curricular activities, and thus it would be physically impossible. Secondly, supposing it existed, such a device would be difficult to adopt if it were special-purpose; school logistics and resource limits constrain classroom technology to those which can support many learning scenarios. Consideration of these practical obstacles has informed many of the decisions in this dissertation.

Inspired in part by the vision portrayed in *Harold and the Purple Crayon*, Crockett Johnson's magical 1950s book about a small boy who can draw objects and scenes that come to life [103], we envision a tool that enables learners to, like Harold, be able to diagram a physical system and then explore it (Figure 1.1).

1.1 Design Considerations

There exist multi-purpose commercial haptic displays such as Phantom (now Geomagic [1]), Novint Falcon (Novint Technologies, 2010; discontinued but currently produced by a different vendor, 2019 [87]), and Force Dimension Omega displays (Force Dimension, 2017 [61]) that suit many applications; however, these devices are often expensive and not accessible in educational settings even as educational toys. Knowing the design considerations tailored towards a specific application enables a more targeted design of haptic displays and consequently reduces the cost. Here, we present a list of design considerations that we collected based on literature and our group's previous experiences in designing an educational haptic device.

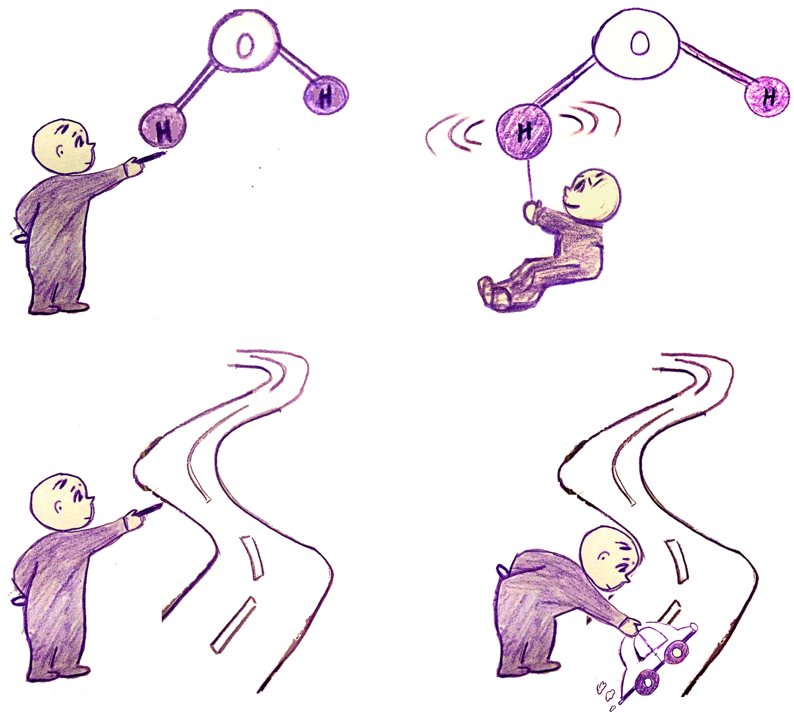

Figure 1.1: Harold can now feel (not just see) what he draws,from molecular attractions (top) to roadside edges (bottom)

1.1.1 Pedagogical Considerations

As we sought to understand how to extend the benefits of digital manipulatives to learning of more advanced topics, we found traction in two visions based on early works by Seymour Papert and his colleagues at MIT Media lab. His two visions include:

- aiding learners to expeditiously express their thoughts by facilitating communication of ideas between human/computer and human/human,
- supporting exploration of different domains of knowledge such as building a model of the world or abstract concepts by making them more accessible or helping to make them approachable in new ways, therefore increasing the possibility that learners use this medium to perform interiorized actions on

> the object of knowledge [180].

Thus, we set out to implement these two visions in the form of a digital manipulative to function as an "object-to-think-with" [174]. We clarify what we mean by an "object-to-think-with" by giving an example of a girl interacting with an educational robot called *Curlybot* [64].

> *"... It usually was not possible to have children perform specific tasks given the informal environment of the study. However, there was one seven-year-old girl that played with Curlybot for an extended period of time and accepted our challenge to create a few geometric shapes out of their most basic elements. We found that she needed us to provide an example before being able to create the shapes on her own. We showed her how to create a square and let her try it on her own. When we asked her to create a circle, she started by designing it with large arcs. She needed additional help to understand that a circle could be created from a small segment. Later on, the same girl came back, and asked if she could try a shape she had been thinking about. We were pleased to see that she continued to process her new knowledge about shapes even outside the play area. Curlybot appears to have become an object-to-think-with for her."*

This example highlights the learner's experience model from becoming aware of how to use the robot to taking actions and further maintaining these actions even without having access to the robot. What Curlybot does is to open up new opportunities by giving the means to learner to approach new problems based on their own experiences. We will explain these two visions in details in Chapter 6.

1.1.2 School Logistics considerations

Ozgur et al [171] propose a list of requirements for a useful educational platform in a class setting. From a practical standpoint, the platform should be low-cost, robust, and reliable. The limited class-hour (typically around 45-50 minutes) requires an educational tool that can support uninterrupted learning scenarios, minimum calibration and initialization time, and at the same time shows encouraging learning gain to justify the effort to design a learning activity with the tool. As such, we should carefully consider the user experience for both learners and the teacher.

Other considerations drawn from some of our own experiences with classroom work include extreme limits on teacher time (both for studies and later deployment) and technological expertise or the ability to prioritize such time as part of their job. On the technology side, we need to be concerned about justifying the expense of sole-use technology, deployment practicalities like batteries and power cords, and the sheer difficulty of students being able to determine when a device is behaving correctly, yet delivering encouraging learning gains and engagement right out of the gate in order to justify the extra workload [170].

This highlights a larger issue, that of validating learning benefit in situations where there are countless variables and controlled studies are not possible. For this reason, many studies take a qualitative approach and look for ways in which the haptic modality is changing student strategies, collaboration style, engagement and interest or type of questions [44].

1.1.3 Challenges of Haptic Displays

Making meaningful haptic feedback is challenging. Often vibrotactile haptic feedback falls short in making the desired haptic perceptions. As a result, haptic designers employ more sophisticated force feedback haptic displays. On the other hand, force feedback haptic devices have traditionally been costly and difficult to use; therefore, inaccessible for many applications. Educational haptic devices were not an exception. Below, we give an example of educational haptic challenges in practice.

An education researcher asked our group for a large workspace haptic device to study the role of bodily gestures in understanding salient features of graphs.

We could not offer them any solution at the time, as haptic devices were not ready to be employed in real classroom settings due to hardware and software limitations. Most haptic devices need to be anchored to a base in order to transfer reaction forces to the ground via mechanical chains of links (arms) or cables. This static grounding constricts mobility and restricts them to a small working area. Increasing the workspace requires expensive and powerful motors and sturdy linkages. Consequently, our options were either to invest in expensive devices or be limited by the need for grounding.

1.1.4 Creating Haptically Augmented Learning Environments

There are a multitude of haptic libraries to support designers developing haptic interactions with virtual environment for a given haptic technology. However, designing even a simple environment using a 2D haptic library requires some basic knowledge of programming as well as of Physics. Either or both of teacher and student may not possess this knowledge or be confident enough to write programs in a classroom setting. Even in cases where a teacher is a technology enthusiast, the uncertainty in predicting learning benefit relative to a large time investment is an understandable barrier. This underscores a broad need for more usable, accessible tools for haptic experience *design* which goes well beyond the need for accessible technology itself.

Roadmap: We will explain the challenges of haptic displays for education in Chapter 2, and our effort to design a new mechanism to address these challenges in Chapter 3. We will present our two studies on educational applications of our new device in Chapters 4 and 5. In Chapter 6, we will discuss our method of creating haptic experience just by drawing it. Finally, in Chapter 7, we summarize the results and provide a conclusion as well as directions for future research.

1.2 Thesis Direction, Rationale and Scope

1.2.1 Thesis Focus and Scope

The prior consideration i.e. the lack of appropriate haptic devices discussed above led to a focus in this PhD research on making a new device that can provide force feedback to a user's pen and paper interaction. We study and define a framework for this type of interaction and investigate how well it could work in terms of the mentioned considerations. We explore the best strategies for employing force feedback in learning activities using three educational haptic devices including the one that we have invented. The use of multiple devices makes the results generalizable and reduces the risks of failure due to the varying qualities of haptic rendering per each device. This dissertation does not cover any studies on potential learning efficacy and any evaluation of whether this hardware approach will be more suitable in a classroom context.

1.2.2 Design Approach

From the wide range of digital manipulatives for education, we decided to explore an intersection of two main streams of research on (a) haptic devices and (b) pen and paper-based interactions with a primary focus on STEM applications. Although there exist several examples of devices that are specifically designed for each mentioned domain of research, only a few of them support both haptic features and pen-and-paper interaction simultaneously.

Our options were either to invest in high quality but expensive haptic displays that we could implement and test our ideas knowing the drawbacks that the small work-space could impose on our research findings, or to design our haptic display based on the design considerations that we have previously mentioned. Because of my engineering background, we decided to go with the second choice. Therefore, a major activity throughout this dissertation research was to to develop the hardware and improve its capabilities.

We decided to make 2D haptic interfaces that can support generic interaction with content found in typical curricula, conceptual learning, and ultimately, amenability to commercial development. The ability to create and feel virtual constraints while generating smooth curves and straight lines on life-sized 2D surfaces could make them more efficient and natural to use. Or, like Harold, a learner could diagram a physical system and then explore it. We believe that it is necessary that learners build the environments they will feel. Current educative haptic platforms must be programmed by hapticians rather than learners. This can result in a passive experimental environment as opposed to creation-based premises of digital manipulatives.

1.2.3 Evaluation Approach

We evaluated how a user can express their ideas more fluidly. We established the core interactions that physically assisted sketching can help users to project out their ideas. As part of our assessment, we measured how force feedback can improve the accuracy of drawing. Moreover, we investigated how a user manages their authority in drawing by giving them control over the amount of aid they receive from the computer.

To test the benefits of haptics for exploring the learning environments which we developed, we used three different haptic devices – the one we developed within this dissertation's scope as well as two others, together exhibiting a diverse set of characteristics – each paired with a learning environment. We studied the benefits of haptics in a collaborative learning context. Therefore, we could explore what strategies learners would take to use this tool to achieve mutual understanding which eventually shed light on the best methods to use haptics in learning scenarios.

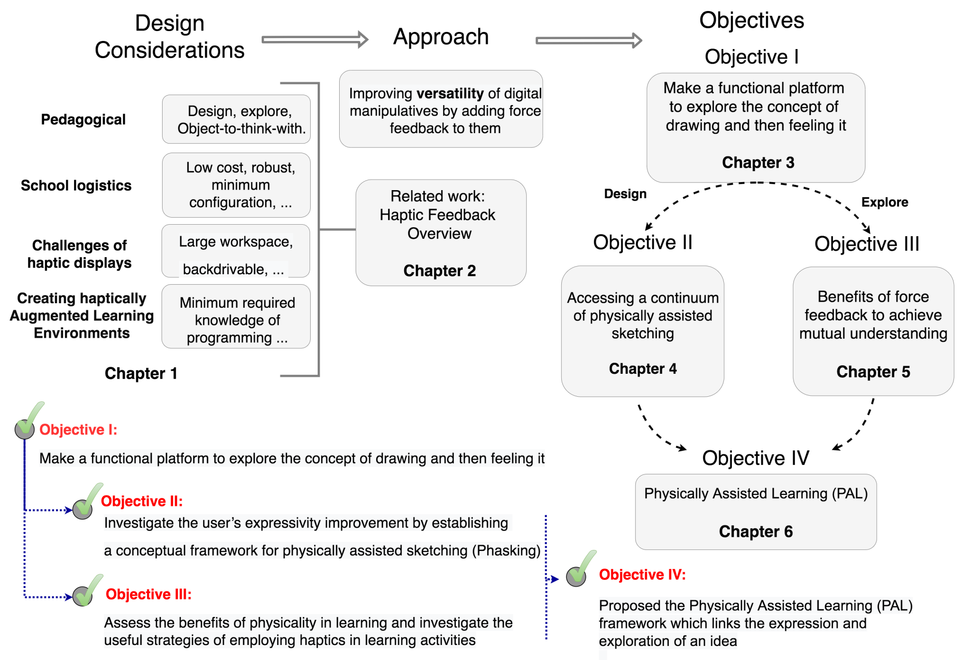

Figure 1.2: Ph.D. big picture overview

1.3 Thesis Objectives

In this section, we define the thesis objectives and describe the specific mechanism and framework we used to reach each objective (see Figure 1.2).

1.3.1 Objective I – Chapter 3

We created a functional platform to explore the concept of drawing and then feel it.

- **Guiding question:** How can we create a low-cost, large workspace force feedback device? What type of new interactions can we support with it? What are the potential educational applications of this platform?
- **Approach:** Explore the design considerations including the mechatronics, user experience, and ergonomics requirements to design a pen-based force feedback device and evolve it through a series of iterations.
- **Achievements:** We designed two form factors of a handheld force feedback

device: “stylus-based” and “pen add-on”. We ran mechanical tests to ensure that they meet the primitive haptic device requirements. We started with user studies that required the stylus form factor due to the simplicity of the design compared to the add-on form factor. We assessed the quality of haptic rendering with the “stylus-based” form factor and the drawing with the “pen add-on”.

1.3.2 Objective II – Chapter 4

We investigated how we can help users to fluidly express their ideas. We formalized our inquiry by establishing a conceptual framework for physically assisted sketching (Phasking) which eventually informed our evaluation.

- **Guiding question:** What are the key interaction concepts for physically assisted sketching (phasking), and how can we support them with our force feedback pen?
- **Approach:** Explore how physical assistance from a haptic display can improve the user’s expressivity in pre-defined drawing tasks. We established the framework of physically assisted sketching by identifying a primitive set of interaction concepts. We use this framework as a base for our user evaluation. A user explicitly sets the drawing commands (e.g., drawing a circle) and receives the proper force feedback. Later, we reflected on the drawing performance and the unique opportunities, which our system brings to the student’s physical drawing.
- **Achievements:** We used an untethered force feedback pen to give real-time force feedback. A user could use the command palette to choose among the functions that he/she needs the assist for drawing. We used informative metrics related to accuracy and shared control and authority to asses the performance of the haptic device.

1.3.3 Objective III – Chapter 5

We studied the importance of sense of touch in experiential learning; and uncovered which haptic strategies will be most useful in collaborative learning.

- **Guiding question:** How can we improve the versatility of the device through force feedback- the capacity to express information to users through the addition of haptics?

- **Approach:** Study how learners can utilize force feedback in the context of a collaborative learning framework, in a manner that is generalizable beyond the device we invented and lends insight into the utility of different kinds of haptic capabilities. Within the collaboration, we specifically look at the process by which two learners achieve a mutual understanding of the concept. Studying with learner pairs allows us to leverage existing theories of stages of collaborative learning, and more easily relate our observations to findings using other kinds of learning technology.

- **Achievements:** We deployed a variety of force feedback devices for a variety of science concepts, each paired with different interactive learning environments. We gained insight into the different approaches and strategies that students take and how the platforms supported or hindered them by means of qualitative methods applied to student pairs.

1.3.4 Objective IV – Chapter 6

Here we propose an approach to connect the activities in design (Objective II) and exploration (Objective III) to achieve a smooth transition between these two stages in learning.

- **Guiding question:** What are the key haptic interactions that can support learners throughout different stages of experiential learning cycle?

- **Approach:** Propose a theoretical framework to support physically assisted learning in two stages of experiential learning namely *Active Experimentation* and *Concrete Experience*. Discuss the types of haptic interactions in each stage and explore the need to link these stages together to create a smooth and natural transition between them.

- **Achievements:** We implemented and technically evaluated the key haptic interactions such as rendering a wall and Mass-Spring examples to validate the feasibility and performance of the concept using available haptic and digital pen technologies then proposed a path forward to more generalizable learning interactions in other domains of Physics.

1.4 Contributions

Our educational contribution consists of the design, implementation, and evaluation of a new haptic digital manipulative that serves learners in their design and exploration of learning environments. We achieve this by pursuing the mentioned objectives (I-IV). Besides the educational contribution, we present three thesis-level contributions, which impact the research in areas other than education.

1. We introduce the Ballpoint drive, a mechanism that can provide force feedback through the friction of a rolling ball on a surface. Our approach circumvents the current workspace limitation of 2D haptic displays by offering an infinite 2D workspace. This contribution is directed towards the haptics engineering community **Contribution I**.

2. We demonstrate how haptics guides users in their physical drawing. We create a platform that supports computer-assisted design in physical sketches as well as digital twinning where it is possible to trace the physical content on the screen and being able to apply changes and bring them back on the paper through guided drawing. Our platform also creates a medium for smooth collaboration between human and computer. This contribution is towards the design community **Contribution II**.

3. We study the added value of haptics in collaborative learning. We find values in assessing the use of haptics in a collaborative context and investigate how different forms of force feedback impact the dynamics of collaboration. We identify successful strategies of using haptics in peers' collaboration. This contribution informs the research in both fields of educational haptics and collaborative learning **Contribution III**.

Chapter 2

Background on Educational Haptic Technology

Preface – *In Chapter 1, we discussed the challenges of haptic displays as well as difficulties on the way of designers who are creating haptically augmented learning environments. An overview of important topics in the field of haptics will help us to better understand these challenges. In this chapter, we provide the relevant technical background for this dissertation. We begin with two classes of haptic feedback and focus on the force feedback's important terminologies as it is the pivotal point of this dissertation. We explain the haptic performance metrics as the basis for the validation of any introduced haptic display. Next, we discuss haptic rendering techniques to enhance sensory experience and available haptic libraries. Having this background enables us to discuss the challenges of haptic displays more in-depth and the related research which try solve these challenges. We close this chapter with a survey on educational haptic displays for designing and exploring.*

To avoid repetition, this overview covers the haptic technology related work common to Chapters 4-6. The related work sections of each of these chapters will focus on material relevant to the specific educational application or topic it covers.

2.1 Haptic Feedback Overview

At the most general level, haptic feedback can take two forms: tactile (skin sensations like vibration, pressure and temperature); and forces (generally called *force feedback*, *kinesthetic display*, *proprioceptive feedback*, etc.) [107]. Tactile feedback can be competent in rendering textures, temperatures and material properties [33, 120], but is less able to display interaction forces, *e.g.*, compliance, inertia and mass, and mechanism dynamics [169]. Thus, when these properties are important, force feedback rendering is often more effective (informative, believable and functional in task completion) than tactile feedback can be.

2.1.1 Grounding and Tethering

Force feedback requires something to push against (*i.e.*, a reactive ground), and thus is easiest to provide as a world-grounded robot arm (*e.g.*, mounted on a desk) [6, 29, 150, 210]. This is known as *grounded force feedback*.

Another concept relevant to haptic feedback and this invention is *tethering*. Tethering is in essence a constraint on mobility, as a consequence of either wires and cords (*e.g.*, to supply power, data communication), or mechanical linkages (*e.g.*, for the user's body to interact with a grounded device through a mechanism with a workspace that is small relative to the desired environment space). Thus, tethering of a device to a restricted anchor point can be a direct consequence of grounding, but it can also arise from other practical constraints. Devices which are untethered typically are freed of one or more of these constraints by one or more of the following approaches: (a) being ungrounded (*i.e.*, unable to provide force feedback); (b) utilizing a body-mounted force ground; (c) using a body-worn battery for power; (d) using either wireless communication for data connection to a computer controller, or carrying the computer on the body (usually the computer is also wirelessly connected to external infrastructure in this case).

However, to be fully untethered, the device must in some manner be free of *all* tethers needed for power, data and force transmission.

2.1.2 Workspace Scaling Factors

As described above, all force feedback devices require a physical ground of some kind. Most world-grounded devices must be anchored to a base in order to transfer reaction forces to a ground other than the user's own body, generally via mechanical chains of links (arms) or cables. To achieve an arm-scale workspace, the force feedback device generally must use large motors, sturdy linkages and high-resolution sensors, all of which increases the weight, size and cost of the device.

Moreover, the device's minimum achievable impedance – an important haptic performance measure which relates to the kind of environments it is able to render – depends on mechanical properties of the device such as masses in the structure, friction, and inertia of the gearing system and actuator. As workspace size goes up, impedance range (the difference between the minimum and maximum impedance that the device can render) becomes smaller for two reasons. First, the minimum impedance that can be rendered will *increase* due to the effect of inertia on links and actuators. Second, the maximum renderable impedance will *decrease* as a result of larger position errors caused by longer and more compliant linkages [15][239].

Thus, it is difficult to make large-workspace force feedback devices using the conventional method of world-grounding, as is the typical approach when systems must provide reactive forces to the user.

2.1.3 Haptic Rendering

Force feedback has long been shown to enhance sensory experience and cognitive information when used to provide correspondent force information during interaction with a virtual environment [158, 169, 214]. The basic approach for spatial force feedback in a virtual model (one way in which force feedback can be utilized) is for the user to move an end-effector of some kind (*e.g.*, the tip of a held stylus, or a bare finger sliding on a surface) while its position is sensed. A reaction force consistent with that spatial point in a virtual model is computed and applied to the user's body part that is connected (directly or indirectly) with the virtual surface [169]. One important kind of virtual feature that is useful to render haptically is a *virtual fixture*: a line or constraint in a virtual model which the user should not be able to cross, but can use as a guide [194].

2.1.4 Open-Source Haptic Libraries

Haptic libraries cover at least two main categories of functions: rendering haptic behaviors, and connecting the haptic interface modality to other parts of the system and experience, be it an underlying virtual model, graphics and/or sound engines and display, managing other forms of user input and control, and in some cases interaction over a network and with other users and entities. While some have been associated with a specific product, most attempt to generalize support for at least a significant class of devices (*e.g.*, CHAI3d [41], hAPI [66]).

Some haptics libraries support advanced rendering of complex deformation and collisions both haptically and graphically for sophisticated environments such as surgical training simulations. For educational contexts, we often do not need such complexity. In contrast, for student-oriented online physics learning materials it is common to see the physical behaviour of an object presented with simplicity via an open body diagram and illustration of applied forces (*e.g.*, [178]).

On gaming platforms, developers use graphics engines to simulate the behaviour of rigid bodies in a virtual world in procedural animations which move realistically and can be interactive within a virtual game world. Hapticians have exploited game engines for their virtual environment modeling, getting graphic display for free and driving haptic output from the VE simulation; this obviatees the need to make or access another physics library for haptic rendering. For example, the hAPI uses a wrapper around the 2D physics simulation library Fisica and turns it into a haptic engine system for educational purposes [66].

2.1.5 Haptic Performance Metrics

An early article on haptic displays for remote interaction with physical or virtual environments suggests three key qualitative requirements for a haptic display, which are still generally accepted as both important and challenging to achieve, particularly while meeting other requirements like low cost, increased workspace size or portability [150]. They can be expressed as primary performance metrics:

1. *Free space must feel free.* This freeness is obtained by minimizing the overall electromechanical impedance when actuators are turned off, *i.e.*, portraying free motion. Impedance can come from many sources, including actuator

non-backdrivability, and friction or backlash throughout the drive train. For example, the easiest way to achieve high forces is through a large gear ratio, but this leads to high motor inertia when the device is backdriven, violating this requirement.

2. *Solid objects must feel stiff.* The rendering of walls and other solid objects requires both at least moderately large forces, and rendering stability at a point of nonlinearity (the crossing of a virtual line between free space and a hard virtual surface). The latter requires high temporal and spatial resolution relative to the speed of the end-effector's movement. Beyond this, to be considered of high fidelity, a device must to be able to faithfully render different levels of stiffness between none and high.

3. *Virtual constraints should not be easily saturated.* Apart from the portrayal of gradations of stiffness (previous item), the issue of saturation is about total force magnitude. When the user hits a virtual wall or slides along it with a reasonable interaction pressure, the device must be able to respond with sufficient force that the user will not overcome it.

To compare haptic displays, we can consider quantified parameters of the above general statements (*i.e.*, actual free-space impedance, rendered stiffness range and force saturation), but also secondary performance measures. The latter include, for example, peak force, resolution and precision, workspace size, degrees of freedom, bandwidth, rendering impedance bandwidth, power density, and peak acceleration [89, 201].

2.1.6 Sharing Haptic Control Authority

Early studies have suggested that haptic shared control can benefit the performance (speed and accuracy) of human and robot collaborative tasks, as well as lowering the need for visual involvement and level of control effort [2, 3]. This seamless collaboration helps users to take the full authority of control, or conversely, shift it smoothly to the system.

In an automation application, haptic displays can be incorporated into the manual control displays, when there is a need for machine-human shared control collaboration. Users can receive relevant information continuously through their sense of touch, and decide to either conform, or override to gain control [77].

2.1.7 Dimensionality: One-, Two- and Three-Dimensional Rendering

A great deal of effort has been devoted to the creation of force feedback devices that can render forces in three dimensions (3D); *e.g.*, as seen in http:/www.haptipedia.org. The need for grounding constricts mobility; devices with large workspaces are complex and costly. Wearable devices capable of rendering forces to, *e.g.*, the hands and fingers in body-referent frames can be untethered, with a less constrained workspace. However, force feedback variants are currently cumbersome, complex devices that provide grounding by pushing against another part of the user's body. Tactile (*e.g.*, vibration) feedback can be easily provided in an ungrounded, untethered form, but is limited in what it can render [213].

Two-dimensional (2D) haptics typically provide feedback in a plane. This has been done in many ways, with both force feedback and tactile display; and with both a held or touched mechanism (*e.g.*, a pantograph mechanism driven by two actuators [186]) or with direct touch (bare finger sliding on a surface that is actuated in some way). Examples of various approaches that have been taken to achieve 2D haptics are described below, together with brief examples.

Single degree-of-freedom (1D) devices can be potentially low cost and simple [72, 146, 149], and are well suited to many interface specifications. However, they are limited in the types of systems they can represent. When built at extremely low cost, as with many other technologies, rendering fidelity as well as workspace size may suffer.

2.1.8 Passive versus Active Haptic Feedback

Energetically *passive* haptic devices cannot supply energy to a haptic interaction, but can only dissipate energy that is supplied by the user, *e.g.*, through the user's own movement. A brake that is under electronic control is an example of a passive haptic display: it can stop or slow the user's movement, but cannot actively guide the user along a trajectory, or simulate storage and release of energy in an elastic material. This characteristic limits the scope of what a passive display can render, but its passivity confers advantages of guaranteed stability.

Active haptic feedback devices, on the other hand, can render a broad range of haptic features, and can also supply active guidance. In guidance, the user's hand can relax while guidance forces draws it along a trajectory or to a specific point of interest.

2.2 2D Technologies with Energetically Passive (Non-Guiding) Mechanisms

2.2.1 Passive Collaborative Robots (Cobots)

In 1996, Colgate *et al.* introduced Collaborative robots *(Cobots)*. Cobots are energetically passive devices, meaning they are unable to apply forces to the user [40, 228]. They operate within a multi-wheel (*e.g.*, tricycle) mechanism, which can be moved around under operator control with no resistance (an actively steered, but passively driven single wheel design). Then, to render or enforce a constraint (*e.g.*, a virtual wall), the appropriate wheel turns sideways and blocks movement in that direction while allowing continued motion in other directions. This class of devices has been used to safely provide constraints in material and parts handling in automobile assembly. The Cobot mechanism is suitable for passively presenting smooth, hard virtual surfaces, and can be integrated with path planning ability. However, it is less suitable for other haptic effects such as texture, friction and compliance. Finally, as a passive device, is unable to provide active guidance forces.

Twenty years afterwards in 2016, Price and Sup developed *Haptic Robot* [183] as a hand-size passive force feedback device with a mouse form-factor, which exploits a Cobot unicycle mechanism (motor drives a caster to change the axis of rotation) to create haptic force feedback. In essence, this device takes the Cobot robot idea, which was designed for materials handling in manufacturing environments, and configures it as a haptic feedback instrument. In this mechanism, a wall collision is rendered when the motor puts the caster in a path that can only steer tangent to the wall. The device can present definitive constraints (resistive force of more than 10 N; but like the Cobot, it cannot create active force feedback or guide users.

2.2.2 Brake-Based Haptic Styli

Invented in 1888, the ballpoint pen is still a ubiquitous writing implement due to its simple but functional mechanism. A number of later efforts, illustrated in the patent literature, have found ways to actuate or brake the rolling ball, albeit generally at a far larger scale and for different purposes. For example many following patents suggest new ideas of how to control the speed of the rotating ball by means of electromagnetic or mechanical brakes. They are all dissipative in nature, *i.e.*, energetically passive.

For instance, **US Pat. No 7,508,382 B2** [128] presents a force feedback stylus with an electromagnetic actuator and rotating ball. **US Pat. No 7,265,750 B2** [195] describes a configuration where a solenoid is embedded inside of a stylus which moves to change the pressure on the rotating ball, preventing it from free rotation; this is essentially a force feedback interface using a mechanical brake to hinder the free rotation of the contact ball. The ball is placed between two supports, while a coiled actuator moves the braking pad and housing to increase the resistance to the motion of the ball under electronic command. A version of this system with a *haptuator* (an actuator specially designed for haptic applications) was designed to render roughness on 3D objects [43].

There exist other haptic styli with dynamic resistance which use different mechanisms to control the friction and the speed of the rolling contact ball; *e.g.*, **US 8,681,130 B2** [7] and **US 8,619,064 B2** [121]. **US Pat. No 9,116,560 B1** [81] describes a touch pen with a similar braking mechanism, but uses multiple haptic balls to improve haptic feedback and accuracy.

None of these friction-based mechanisms are designed to provide active force feedback to the user: specifically, there is no mechanical driver to roll the contact ball.

2.2.3 Other Haptically Enabled but Passive Styli

There are examples of haptic styli that use friction modulation to provide a slightly rough pen-on-paper sensation on glassy display surfaces. Advantages include greater controllability and a lower pressure requirement while writing. The friction between stylus and surface depends only on the user-applied pressure and can be controlled externally [132].

I-draw, 2014 [59] presents the idea of adding an attachment to any stylus-shape object to add force feedback. It suggests a Cobot-like mechanism using a small caster (controlled by a motor) and three rolling balls. This design has the usual Cobot advantages and limitations, and at the same time, the challenge of dexterity. Using one wheel, I-draw can only resist movements perpendicular to its wheel. As a result, it only renders 1-D passive constraints at a time. Eventually, it faces difficulties in rendering sharp corners and can't stop the user's movements in 2-D.

2.2.4 Active Force Feedback Technologies

Active Force Feedback through Robotic Arms: Serial and Parallel Linkages

Robotic arms are among the most known and studied active force feedback haptic devices. The commercially available Phantom (now Geomagic [1]), Novint Falcon (Novint Technologies, 2010; discontinued but currently produced by a different vendor, 2019 [87]); Force Dimension Omega displays (Force Dimension, 2017 [61]) have been used in many different areas including medicine and education. Many studies have investigated ways in which haptic feedback can improve learning processes, for example in teaching visually impaired people how to write [164] or providing more realistic training environments for medical practitioners [37]. Serial arms, often tendon-driven as in the Phantom, can potentially provide relatively large workspaces, but at the cost of precision, strength and cost. Parallel robot mechanisms such as the Novint Falcon can provide great dexterity in their entire working space and tend to be more stiff (allowing higher fidelity rendering); however, their working space is often much smaller than serial robots.

One of the most common 2D force feedback mechanisms is the *Pantograph*, 1994 [186], a 5-bar linkage . These are capable of high-quality feedback in a planar workspace, generally of fairly small size. Recently, a low-cost version was designed to be clipped on top of a tablet display to be used in co-location with the tablet's graphics [67]. Pantographs are appropriate as small-workspace desktop devices but not suitable for mobile applications in currently-available configurations.
Active force Feedback joystick-trackball: In the joystick (released commercially as

the Microsoft Sidewinder consumer device), two motors are connected to a 2DF stick through an advance belt gear mechanism, an approach which is fundamentally different from that of the proposed invention. For the haptic trackball, the mechanical control over the rotating ball is not explained, but one example can be found here [54]. However, for these proposed configurations, the mechanism is sufficiently bulky that it requires the device to be held against a ground surface, such as a table.

Active Force Feedback through Magnetic Forces

Magnetic force can be used to create directional force feedback. The (*FingerFlux*, 2011 [224]) device provides attraction, repulsion, vibration haptic feedback on a planar (2D) tabletop using an array of coil magnets embedded in the table and a permanent magnet mounted on the user's fingertip. A permanent magnet attached to the user's finger responds to an actively controlled magnetic field generated by electromagnetic actuators on or under a tabletop. The results showed reduction of drifting, finger guidance of the finger and physical constraints. However, this approach has practical limitations. Electromagnetic coils are bulky and heavy, presenting challenges to inclusion in planar displays. It is also difficult to achieve a high-resolution magnetic field and control the direction and magnitude of the force. Finally, resultant magnetic fields can be a major source of electrical noise that impacts on other system components.

2.3 Devices Rendering Texture and Friction

Vibrotactile surface display: The most common and broadly commercialized surface-rendered haptic feedback is vibrotactile, produced by high-frequency vibration of a touched surface in response to finger motion. A variety of actuator technologies are used to produce these vibrations. No lateral resistance to motion is provided, *i.e.*, this is not a force feedback technology.
Variable friction: Surface friction is felt in sliding contact of a bare finger with a surface. Some technologies vary the coefficient of friction as a function of measured finger position, and thereby control linear and shear forces experienced by the finger during active sliding. Significantly, variable friction approaches are not able to generate any forces on a stationary finger, but they can generate small amounts of

force resistance to a moving (sliding) touch. Applications enabled on touchscreen devices with this kind of feedback have been explored [134].

The (*TPaD*, 2011 [227]) device use an air squeeze film effect to lower friction. Specifically, it uses ultrasound vibration to create a thin layer of air between finger and display and lower the friction at a designated point. The (*ShiverPaD*, 2010 [34]) employs a similar variable friction mechanism in addition to lateral vibration to create an active forcing device. Despite its innovative mechanism, *ShiverPaD*'s current design suffers from a small workspace (15 mm in diameter) which limits the finger exploration area. Increasing the workspace size demands heavy engineering investment and expertise.

In a different approach to providing variable friction, *TeslaTouch* uses electrovibration principles to create a dynamic friction between the sliding finger and display [16]. Increasing the charge density on a conductive transparent electrode underneath the display can create an electric field between finger and electrode. The alternating electric field can generate periodic electrostatic forces. These forces are too weak to be perceived in a static environment but they can actuate the skin and create a rubbery sensation when fingers are sliding on the surface of the display. Practical levels of electrostatic friction display require a large electrical field, from a high voltage difference between user's finger and touch display. This can potentially be a safety hazard, and requires some mitigation measures to avoid electrical shocks, but this issue has been handled to some degree in products using this approach that are close to release today (*e.g.*, https:/www.tanvas.co). However, still more needs to be done to reduce the required volume for incorporating this technology into commercial tablets and surfaces.

Magnetic texture: Magnetic coil arrays were investigated to create virtual surface roughness directly to the fingertip equipped with a magnet [20]. A user needs to put a "haptic probe" on their fingertip, or in another configuration, hold a pen with embedded coils inside to receive the force feedback from the screen.

Ferrofluid and magnetorheological (MR) fluids: Ferrofluid materials are colloidal liquids made of nanoparticles of ferromagnetic materials that can change their viscosity as being exposed to a magnetic field. Changing this viscosity locally can create labyrinth feeling when the finger passing it. This phenomena was first explored as a surface display in the early 1990s [160]. More recently, *MudPad* is a multi-touch screen based on MR fluid combined with an array of electromagnets that can change the surface stiffness and render changeable textures [102].

One problem with ferrofluid materials is the weak generated resistive force, which improves as the size of the particle increases. Another issue is due to the large sizes of ferromagnetic particles in MR fluids; ferroparticles will settle out if the device remains static for a period of time. An approach currently under investigation is to increase sedimentation time by means of surfactants.

2.4 Challenges of Haptic Displays

2.4.1 The Challenge of Grounded Force Feedback

Traditional grounded force feedback: We will not attempt a comprehensive review of this large and diverse class, but examine two archetypal examples which capture the constraints we wish to circumvent. The 3D Phantom Omni (now Geomagic Touch, geomagic.com) desktop system has been used as a handwriting training tool and for rehabilitation [164]. In 2D, the best known is a pantograph (first in 1992 [141], recently updated in DIY form [67]).

The Omni has a serial cable-drive mechanism; the pantograph uses parallel links. While both mechanisms' workspace can be scaled up, practical considerations limit it to inches per side (about 5" for the Omni). Beyond this requires significant cost to maintain performance.

Force feedback grounded within a surface: (dePENd, 2013 [231]) uses actuated magnets moving under a table to apply directional forces to pen and paper interaction, exploiting the ferromagnetic property of pen ball tips to assist sketching. While it scales to a larger workspace than some devices, it has a large, fixed footprint. It can provide position but not force control, and requires calibration when the user lifts the pen off the surface (a problem we must also creatively address with our ballpoint drive).

Ungrounded tactile feedback devices: at least two reported haptic styli provide *tactile* feedback related to position on a graphically rendered surface, *e.g.*, using a friction brake [43, 91], which allows the device to render precepts like roughness. This dissipative approach, however, is not able to generate directional force feedback or guidance forces, and hence is unable to represent constructs such as springs or walls.

2.4.2 The Challenge of Mobility and Large Workspace

One approach to growing a workspace is to move the ground.

Movement without forces: (Ballbot, 2006 [129]) is a mobile robot based on an inverted mouse drive. While not designed for force rendering, its locomotion mechanism has high relevance to our ballpoint drive. Making this approach suitable for haptic rendering, however, requires adaptation of the force display mechanism, an entire new family of rendering techniques and control strategies, and multimodal integration. We will return to Ballbot as we introduce our own design.

Movement with forces: There are two prior examples of mobile robots (*i.e.*, bots which can propel themselves on a surface) being used as *active* force feedback displays. While perhaps suited to the applications for which they were designed, each has traits which make it unsuitable for high mobility, nomadic applications.

MOTORE (MObile roboT for upper limb neurOrtho REhabilitation, 2011 [12]) is a rehabilitation platform, *e.g.*, re-learning motor skills. Designed to move across a table-like surface, its three wheels allows the device to generate spatial haptic feedback which a user experiences while moving the device around on a surface. It can control and constraint impedance and help patients tracking a trajectory in addition to simulate virtual environment with different viscosity. However, it is 10kg with a 145mm radius, and not comfortably lifted or held

(*Cellulo*, 2017 [171]) is a mobile haptic robot is meant for classroom learning, and able to render virtual objects on a 2D-plane. Cellulo is just 168g, the weight of a pear (although physically larger); and it is low-cost at €125. It uses a permanent magnet ball in its omnidirectional driving mechanism to address known vibration issues for this locomotion approach. Using permanent magnet coupling between motors and the rolling balls provides Cellulo with partial backdrivability.

Cellulo is the sole example we have found of a grounded, mobile surface force display with handheld potential. Its puck-like form affords an enclosing grip and a mouse-like interaction style. However, while inspired by a "pen and paper" metaphor, it is awkward for drawing; and its size produces occlusion issues. Further miniaturization might be problematic for stability reasons because its multiple-ball surface contact relies on a flat orientation.

2.4.3 The Challenge of Accessibility

The "accessible haptics" movement: Large-workspace, low-cost devices will be valuable in many applications; but even with prevalent DIY fabrication strategies like 3D printing and laser cutting, requirements for structural strength, precision, motor quality, and power systems drive prices up exponentially with workspace area. Promisingly, recent haptic prototypes suggest that affordable devices are in fact becoming feasible [145] [149] [62] [67] [176].

Applications needing access: We are engaged with the goal of using haptic feedback to assist with conceptual understanding. While it is beyond our present scope to share the theoretical underpinnings of this goal, some recent studies suggest both its promise and its challenges [44, 155]. Meanwhile, there is evidence that large movements better support embodiment in learning than small ones [69].

2.5 Haptics for Designing and Exploring

In this survey of education-related haptics, we focus on the intersection of two primary haptic approaches: (1) haptically rendered virtual environments, and (2) pen-and paper-based interactions (Figure 2.1).

2.5.1 Design Approaches: Input Methods, Feedback Modalities and CAD features)

We identified a number of relevant works describing novel input methods and haptic feedback outputs; we focus on systems that are well-suited for educational applications such as science, technology, engineering, and mathematics (STEM) learning and visual art.

SketchPad [212] was a pioneer in the field of modern computer-aided design (CAD) as well as pen interactions with graphical displays. Sketchpad, for the first time, demonstrated the great potentials of combining computing power and digital design in 1963.

VoicePen, 2007 [88] is a digital stylus that takes non-linguistic vocal inputs

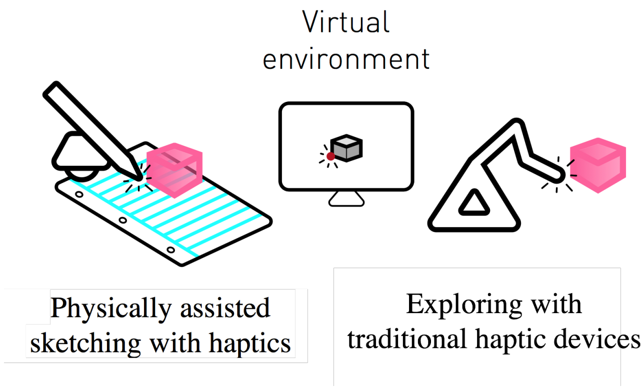

Figure 2.1: Haptic interaction with virtual environment. Two most studied applications of haptics with potential benefit in education. The first group assists users in their sketching and design while the second group is used for exploration.

in addition to position and pen pressure for tasks such as creative drawing and object manipulation . VoicePen uses vowel sounds, variation of pitch, or control of loudness to generate fluid continuous input to the user's pen interactions. Similarly, *WatchPen* [96] captures the users' drawing inputs form a combination of smart gadgets including a digital stylus, a smartwatch and a tablet. It employs vocal and touch input to reduce workflow interruptions, such as tool selection. These sytems' reliance on the vocal modality makes them impractical for classroom settings, but they deliver ideas for stylus interactions.

TAKO-Pen, 2015 [124], is a haptic pen which provides pseudo-force feedback by creating the sensation of sucking on users' fingers through pressure chambers embedded on the handheld surface. *RealPen*, 2016 [32] is a digital stylus which recreates the sensation of writing on real paper with a pencil through auditory and tactile feedback. *FlexStylus*, 2017 [58] allows users to perform tool selection and to draw with better accuracy by bending the pen in various modes. Although these novel input methods and feedback modalities expand the interaction space between users and haptic devices or digital styli, they are designed for specific purposes, and thus are not pivotal for as more broadly useful sketching tools.

In addition to devices, we looked for innovations in computer-aided drawing

(CAD) features for generating engineering or artistic drawings. Specifically, we have found some existing work on parametric sketching – a CAD functionality that allows users to define geometric entities with parameters, and to specify relationships between them as constraints. Examples of parametric sketching include defining a circle by its central position and radius, and defining two lines as being collinear or of equal length. This functionality is useful to architects and architects for creating complex architectural or mechanical sketches. Gürel *et al.* [80] studied the impacts of applying parametric drawing with CAD tools on creating architectural designs, and they found allowing designers to define parameters and constraints on geometric entities provided them with greater flexibility in their creation process. Ullman *et al.* [218] emphasized the importance of geometric constraint as a functionality for CAD tools to help designers create mechanical sketches with better clarity.

Departing from CAD features and considering direct-sketching input, we highlight *ChalkTalk*, which recognizes users' strokes and translates them into meaningful interactions using dynamic visualization and procedural animation to facilitate exploration and communication [179]. ChalkTalk is a purely visual medium at this time; we see potential for using its approach when extending the functionality of the haptically supported approach described in Chapter 6 for more expansive sketch interpretation.

2.5.2 Pen-based Sketching Tools

Engineering Design and Educational Drawing

Aside from work on design methodologies, there are works that focused on sketching on 2D surfaces alone.

In engineering design, *InSitu* provides architects with a stroke-based sketching interface capable of augmenting sites' contextual information from sensor data into sketches and delivering the information via pop-ups ([173]).

Within educational drawing support for (STEM subjects), most devices were built for sketching math or physics diagrams and equations. *MathPad2* allows users to create animations to represent processes (*e.g.*, a mass block oscillating) in addition to static diagrams or math formulas [130]. *Hands-on Math* places more emphasis on recognizing handwritten math inputs from users and performing calculations such as solving for an unknown variable in an equation [237].

Data Visualization and Digital Annotation

In data visualization, *DataToon* utilizes touch and pen-based interactions to allow users create data comic storyboards [119].

Within digital annotation, *PapierCraft* [138] is a haptic stylus that allows users to annotate documents via pen-based interactions. Moreover, it can take users' gestures as document manipulation commands. A later refinement of Papiercraft employed more feedback modalities for distinct purposes: LED lights for notifying users of pen status, tactile feedback to provide warnings, and speech feedback for indicating action results [137].

In this chapter, we have reviewed several key topics in haptics and the challenges of haptic displays. It is apparent from this review that conventional approaches of producing grounded force feedback require mechanical linkages (*e.g.*, arms or cables) to ground the user's hand to a fixed point in space, and thus have workspaces limited to the reach of the devices' arms or cables. In the next chapter, we propose a novel haptic device that displays world-grounded forces to the user's hand in two dimensions (2D force feedback display), over a workspace of unlimited extent.

Chapter 3

Device Design: Introducing the Ballpoint Drive Mechanism

Preface – *Our first step towards evaluating the concept of drawing and then feeling an idea is to make a functional platform that supports the required haptic interactions. This chapter highlights our innovative approach towards solving the need for combining grounded force feedback with a large drawing workspace by introducing a novel 2D haptic display featuring the ballpoint drive mechanism. It describes the evolution of the design of two versions of MagicPens which will be utilized in later chapters, as well as their technical details supported by primary performance measurements. The first version of our prototype has a stylus form factor, and is used for our study on collaborative grounding in Chapter 5. The second version is modified to have an aligned pen-holder. We will describe how we used it in assisted sketching (Chapter 4). This chapter thus comprises the device design sections of two published papers (Kianzad* et al. *2020* [1] *and Kianzad & MacLean 2018* [2]*). The reminders of the aforementioned papers are covered in Chapters 4 and 5, respectively.*

[1] S. Kianzad, Y. Huang, R. Xiao, K. E. MacLean, "Phasking on Paper: Accessing a Continuum of PHysically Assisted SKetchING. Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, 1–12.

[2] S. Kianzad and K. E. MacLean, "Harold's purple crayon rendered in haptics: Large-stroke, handheld ballpoint force feedback", 2018 IEEE Haptics Symposium (HAPTICS), San Francisco, CA, 2018, pp. 106-111.

(a) A rolling, stylus form factor enables unlimited user-driven or guided motion, on arbitrarily oriented and shaped surfaces. We used this in our collaborative learning study, Chapter 5

(b) For a haptically enabled pen, the embedded camera provides an accurate absolute position on a watermarked paper. We exploited the inking on paper feature for our user study in Chapter 4.

Figure 3.1: MagicPen device form factor.

3.1 Overview

An elusive and audacious vision in haptics is the rendering of forces "in the large" – to feel sensations in large gestures, anywhere and anytime. The reality is that delivering force feedback in informal settings and diverse environments suffers from linked constraints. Producing force requires physical grounding, which limits workspace size; meanwhile, expansion impacts affordability and application versatility.

Inspired by needs for haptic support of large motions on a surface in applications like embodied conceptual learning, commercial design, and 2D virtual / augmented reality, we present the *ballpoint drive*. This novel approach circumvents conventional constraints by imposing a new one: motion restricted to rolling on an arbitrary two dimensional surface, and grounding forces generated through friction. We analyze the ballpoint's design considerations in the context of our framing applications, describe a first prototype and its performance, and assess its potential for further development.

3.2 Introduction

Haptic force feedback first appeared in the 1990s [159], and soon fired the imagination of engineers, scientists and children with the possibilities of making a virtual world tangible. Since that time, device designs have proliferated.

However, the fundamental requirement of *physical grounding* to provide a solid reaction force still dictates how force feedback devices are configured, triggering a cascade of interlocked challenges. Most devices achieve grounding by being fixed to the world, *e.g.*, through a parallel or serial robotic mechanism. This imposes a finite *workspace size* – a serious detraction for many non-desktop applications, particularly those in virtual and augmented reality (VR/AR). Increasing workspace usually means large linkages, more expensive components and complex engineering, and hence a sharply increased cost/performance ratio that limits *accessibility* for many applications. Finally, device bulk, tethering and handle restrict *versatility* in how the end-effector can be grasped and moved. While a few applications can live within these four linked constraints, many more are precluded by one or more of them. The most common solution is to get by without grounded force, for a far less compelling result.

A different tradeoff is possible. **Restrict motion to a ball rolling along a two-dimensional (2D) surface;** then, *grounding* comes from pushing against the surface, while the ball driver provides 2D constraints and active forces. *Workspace* is unlimited, or at least as large as the surface. The device itself must be handheld, small and light.

At least one mechanism can achieve this *accessibly*, with low-cost components and DIY (Do it Yourself) fabrication techniques. Variations of the basic design can support *grip versatility* with "swappable" handles and varied orientation.

We have dubbed this concept **MagicPen**. It is inspired in part by *Harold and the Purple Crayon*, Crockett Johnson's magical 1950s vision of a small boy who can draw objects and scenes that come to life [103]. Our surface might be a large wall-mounted or table graphical display, or a personal tablet. It could also be a whiteboard or an arbitrary large curved object, together with a projected image and vision system to detect drawn graphical symbols as a virtual environment is sketched on the surface.

2D haptic interfaces have broad potential applicability, *particularly when they can be arbitrarily large.* This includes generic interaction with common content, conceptual learning and commercial design. For example, automotive designers use tape drawing to create fast, 1:1 scale sketches, and while digital versions exist [78], the ability to create and feel virtual constraints while generating smooth curves and straight lines on life-sized 2D surfaces could make them more efficient and natural to use. Or, like Harold, a student could diagram a physical system and then explore it (Figure 1.1). Handheld haptics are the most novel component of this vision, and the one we address in this chapter. We contribute:

A novel portable mechanism providing force feedback while rolling on a 2D surface, through interchangeable handle forms, whose cost and performance is insensitive to workspace size;

Associated measurement techniques for 2D position and angular orientation;

Accessible designs for user community development.

3.3 Approach: Design Considerations

We describe initial balancing of perspectives of mechatronics, cost, computation, ergonomics, usability and application needs; specific engineering requirements continue to evolve.

3.3.1 Mechatronics and Amenities

Accepted mechatronic requirements for haptic devices include maximal *backdrivability*, minimal *impedance* in free motion, *force* responsiveness and *sensing* bandwidth [198] [90]. Other specifications are dictated by our approach and applications.

Force magnitude: Forces must be clearly perceptible in guiding and rendering, but need not overpower the hand.

Degrees of freedom: Force vector must be smoothly controllable in a 2 degrees of freedom (DF).

Surface contact: The device must maintain rolling contact with a comfortable grip; *e.g.*, rendering a circle requires 2D force vectors without slip at the virtual object's circumference.

Interaction surface properties: The contact ball must be able to maintain frictional grip and optical sensing quality on surfaces that are to some degree rough, slippery or reflective.

Size and weight: A handheld that works on both vertical and horizontal surfaces requires minimal weight and girth, and suitability for an untethered, battery-powered package.

3.3.2 User Experience and Ergonomics

We start with well-known features of everyday pen use. Pens should be comfortably light to hold, but weight imparts stability, normal force and a sense of quality. Their long aspect ratio requires balance. Pen shape and surface properties influence grasp and consequent fatigue.

Ergonomic design for a haptically active pen is more nuanced. We observed people writing with a variety of everyday writing tools and with mouse motions, as we pulled and pushed on their implement with a string or wire. We found that a pen required less force (~0.3N) to drive the user's hand than did the mouse-shape form factor (~0.8N).

Our vision includes being able to sketch environments for rendering, and thus we examined integration of a writing implement with the ballpoint drive. It is clearly desirable to minimize distance between the implement's tip and the force feedback drive contact, to avoid ambiguity between sensing and contact representation. Some provision can also be made through manipulation of visual-haptic co-location.

Based on these considerations, we focused our design explorations on factors of *shape*, *weight*, and *grip*.

3.4 Prototype Design

3.4.1 Form factors

We developed two different form factors as shown in Figure 3.1. We exploit the features of each form factor for our user studies in Chapters 4, 5. Here we elaborate more on these features.

MagicPen with stylus form factor

Our basic proof-of-concept MagicPen is stylus-based (Figure 3.1a). Through a pen grasp, users feel a directional force generated through rolling contact between the driven mouse ball in the tip, and an arbitrarily shaped or sized 2D surface possessing nominal surface friction. This design can haptically render dynamic 2D virtual objects: *e.g.*, constraints, textures, detents, stiffness, viscosity and inertia.

Size and form factor: The bulkiest element is the tip, which houses the ball drive. Its current 45mm diameter is significantly more compact than possible with a mouse form [171]. The stylus grip is familiar and dexterous, facilitating drawing and stroking as well as nuanced perception of guiding forces.

We found that the stylus requires substantively smaller drive forces than a mouse form for ergonomic reasons. The ball drive's freedom to tilt (currently 25-30°) means that the user's applied force is generally not aligned with the surface normal. In a pen grip, users generally take responsibility for maintaining a comfortable tilt angle. This results in their contributing to the manual work of moving the device, with a tendency to follow the pen when it 'pulls out from under' their hand. Further, users tend to rest hand and forearm on a mouse, adding to the inertia that must be overcome with a driver; this happens far less with a stylus. These factors differ most dramatically from a mouse on a horizontal surface, but are somewhat in play in a vertical setting.

MagicPen as Add-on to Arbitrary Pen Implement

Our device can be attached to a pencil, whiteboard marker or pointing finger, and will apply force feedback through the implement while rolling-ball surface contact is maintained. The device barrel unscrews, to be replaced with a retractable pen holder 3.1b.

The substantial increase in application possibilities does bring mechanical complexity. Added girth makes the grip slightly awkward, while the dual points of contact constrain stylus pitch angle. Meanwhile, the offset between drive and implement contact points varies depending on grip pitch, requiring additional sensing and correction.

In this configuration, we have constrained pitch to 45°, respecting users' typical

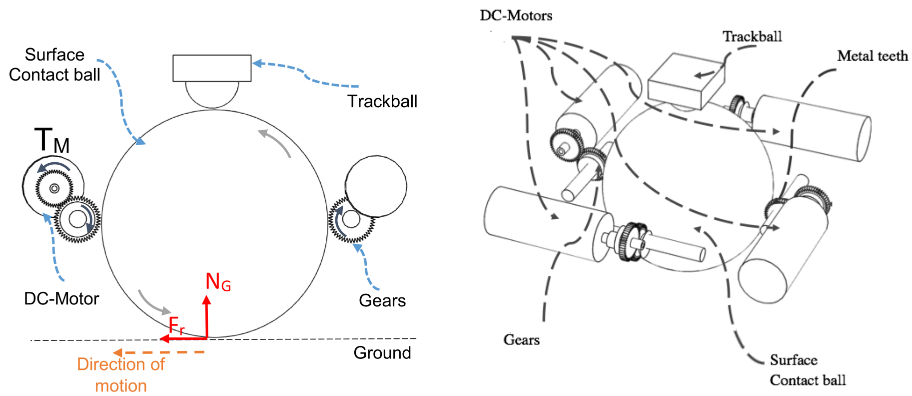

Figure 3.2: Ballpoint drive and sensing in 2D and 3D. **Left:** Motor-generated torques T_m are transmitted to the surface contact ball through the gears causing the ball to roll over the surface. The friction force F_r between the contact ball and the surface produces motion in x-y direction. **Right**: The arrangement of four DC-motors in 3D shows how the Ballpoint drive generates forces in 2 directions.

40-55° writing angle. In practice, users prefer slightly different writing angles on horizontal and vertical surfaces. This issue can be handled with a compliant, angle-sensed linkage with offset corrected in software.

3.4.2 Ballpoint Drive Mechanism

The key driving element of MagicPen is a novel mechanism that we called *Ballpoint drive*. Ballpoint drive (Figure 3.2) is inspired in part by a class of autonomous robots which self-balance on a sphere. Similarly to Ballbot [129], we use four motors (2 opposing pairs) in an inverse mouse-ball drive to generate planar rolling movement (an alternative, three independent drives, leads to nonholonomic constraints). Our innovation is to use this mechanism to render model-based force feedback to a user, and begin to address the new challenges this raises.

Each motor pair acts to rotate a 25.4 mm elastic surface-contact ball, which is retained by a housing with four sprung ball bearings with stiffness tuned to maximize rotary freedom while achieving ball retention.

Components – We tested several motor/gear systems to optimize speed, power and backdrivability. This version employs high power-density brushed DC motors, rated at 2000 W/kg, 50,000 RPM, 3.7V and 100mA. Generally used in small hobby aircraft, they are China-sourced at $2.67 each. In practice, we observe 200-500mA draw and a maximum of 100 W/kg. For context, the high-fidelity haptic standard Maxon RE40 is rated at 300 W/kg for >$300.

The mechanical transmission amplifies torque in two stages: 12:1 from contact ball to rotating bar, and 2.5:1 from the plastic crown gear to the motor. We varied ball material and size to ensure adequate friction on a variety of surfaces. Traditional mouse balls have internal weight which adds excessive inertia. We found best results with spheres of solid neoprene rubber (mcmaster.com/rubber-balls), achieving good contact with monitor screens, whiteboard, painted wall, and tabletop. More extensive research is required to precisely measure actual friction achieved.

Sensing: position, orientation and pressure – The two primary sensing needs for phasking operations are (a) localization of device on the interaction surface, and (b) internal ball motion and stylus orientation for closed-loop control on position and velocity, to generate desired forces.

Position for external localization and internal control:

There are many possible approaches to external localization of the contact point on the interaction surface, depending on application setting, priorities and constraints. We exploited existing technology of digital pens and watermarked papers to obtain absolute position sensing. An embedded camera (here, a Neo SmartPen [98]) decodes optical microdot patterns on watermarked paper to determine absolute position of the stylus to an accuracy of 0.1mm. This enables accurate sketches and useful interactions such as the tool palette.

Internally, a micro-trackball senses the motion of the larger surface contact ball in each rolling direction and sends the pulses to the device controller.

Orientation: Stylus orientation corrects position estimates generated by the above methods. A twist of the user's hand impacts the contact relation between rolling ball and interaction surface: when non-vertical, the ball's contact point does not coincide with the stylus axis. The global localization signal can also suffer from low update rate, gaps or spatial inaccuracy.

We used a low-cost orientation sensor (Bosch BNO055) which fuses accelerometer, magnetometer and gyroscope data to supply Euler roll, pitch and yaw [25].

Pressure: The Neo Smartpen supplies pen tip pressure (0-255), which we use for control sharing. During early prototyping, we tested several locations for pressure sensing. We found out that users have most control over pressure application when the pressure sensor is located at the pen tip, rather than behind the trackball or under the pen finger's grip.

Battery – The early version of the ballpoint drive (Figure 3.1a) was electrically tethered to external power and communications. The next prototype (Figure 3.1b) uses two 3.7V Ultrafire 14500 rechargeable batteries in parallel, capable of approximately one hour of free-roam performance. The most significant power draw is from the CPU due to I/O interrupts related to orientation and localization sensing. A future version using an embedded CPU will be significantly more efficient.

3.4.3 Prototype Software Architecture

Our initial compute structure optimized quick construction over optimal performance. It has two versions of computational elements: an Arduino Mega for sensing, output and kinematics computations to a laptop-hosted custom Python API for the VE model for stylus form factor (Figure 3.1a) or integration of both in a Raspberry Pi zero W (RP) (Figure 3.1b).

The Arduino/RP receives desired force/velocity from the VE model, transforms them to V_x^d, V_y^d based on trackball and orientation sensor data, and generates PWM motor output (U_x, U_y). The laptop API/RP receives kinematic data and both graphically displays and returns VE output.

This system achieves 1ms (1000 Hz) roundtrip updates, limited by the USB connection. An upgraded communication protocol will easily attain the haptic standard of 1 kHz.

3.5 Force Feedback in a Ballpoint Drive

We have framed initial control of our ballpoint haptic display, constrained to a 2D surface, in modes of *guiding* along a path (user is led by the device), or *rendering* a virtual environment (user drives while interacting with VE elements). We implemented two proof-of-concept rendering schemes to test the mechanism's viability in force feedback.

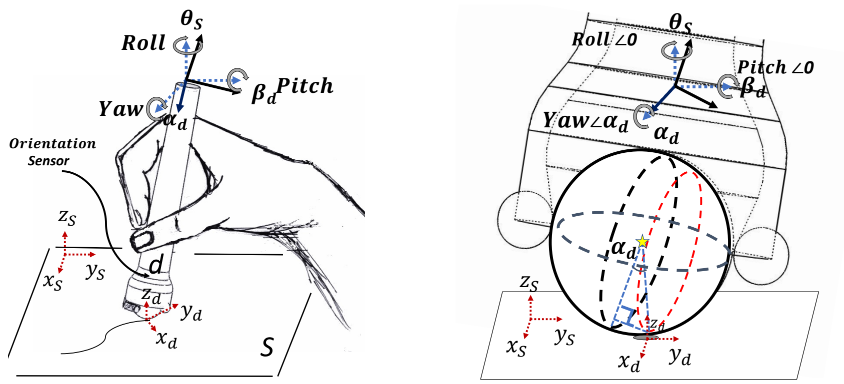

(a) Stylus orientation (dictated by user grip) is defined by roll, pitch and yaw from surface normal. The ball contacts the surface at x_S, y_S.

(b) Surface contact diverges from ball "south pole" when stylus yaws. Eqs. 3.3-3.4 address the resulting parallax.

Figure 3.3: Ballpoint rectilinear and rotational coordinate systems.

3.5.1 Guiding Along a Path: Position and Velocity Trajectory

In this simplest-possible scenario, the device is driven as if a standalone ballbot [129] on a defined trajectory. The user is pulled by the handle without influence on the computed path.

Our ballpoint display can take any pose $[x, y, \Theta_S, \alpha_d, \beta_d]$, where x_S, y_S describe *contact point* in the traversed, topologically contiguous 2D surface's coordinate frame **S**. The remainder are respectively device *roll* around the surface normal, and *pitch* and *yaw* relative to device base **d** (Figure 4.1a).

We do not actuate Θ, α or β, although technically possible with added complexity. Guidance scenarios exist where actuated roll is valuable; for most uses, we believe these are best left passive, allowing the user to find a comfortable angle rather than imposing unnatural torque on the wrist.

The device is driven by commanded forces (F_x^d, F_y^d) applied to the ball, causing it to roll or resist rolling in x_d, y_d. Passed through our voltage-controlled, opposing-pair drivers, these are roughly proportional to the resulting output velocity (V_x^d, V_y^d) when the ballbot moves unresisted by a hand, which are in turn related to surface velocity through the bot's inverse kinematics as a function of roll, pitch and yaw:

$$\left\{ \begin{array}{c} V_x^d \\ V_y^d \end{array} \right\} = R(-\Theta)Y(\alpha)^{-1}P(\beta)^{-1} \left\{ \begin{array}{c} V_x^S \\ V_y^S \end{array} \right\}, \tag{3.1}$$

where:

$$R(\Theta) = \left\{ \begin{array}{cc} \cos(\Theta) & \text{-}\sin(\Theta) \\ \sin(\Theta) & \cos(\Theta) \end{array} \right\} \tag{3.2}$$

$$Y(\alpha) = \left\{ \begin{array}{cc} \cos(\alpha) & 0 \\ 0 & 1 \end{array} \right\} \tag{3.3}$$

$$P(\beta) = \left\{ \begin{array}{cc} 1 & 0 \\ 0 & \cos(\beta) \end{array} \right\}. \tag{3.4}$$

$R(\Theta_S)$ compensates for user wrist rotation (spin), while $Y(\alpha_d)$ $P(\beta_d)$ accounts for grip-generated stylus tilt in yaw and pitch directions respectively.

Figure 4.1d shows that as the stylus tilts sideways in yaw, contact point diverges from the rolling ball's south pole. Now, if the ball rolls toward us (out of the paper plane), it moves through a sub-equatorial circumference, path P_t. This smaller contact circle's diameter is a $cos(\alpha)$ function in x_d and $cos(\beta_d)$ for y_d. We measured stylus orientation with the BNO055 (blended accelerometer, magnetometer and gyroscope data fused in three-axis roll/pitch/yaw Euler orientation output; Adafruit.com) mounted on the handle near the ball, sampling stylus absolute orientation at 100 Hz.

Stylus endpoint velocity at point t on path P_t is:

$$\frac{V_x^S}{V_y^S} = \frac{g_x^P}{g_y^P}, \tag{3.5}$$

where g_x^P, g_y^P define P_t's gradient. We scale velocity components to the maximum supported in either x or y alone (the device could move faster on the diagonal where both drivers are in play). This relation enforces a uniform top velocity regardless of (x, y) heading:

$$\sqrt{(V_x^S)^2 + (V_y^S)^2} = V_{x,max}^S. \tag{3.6}$$

3.5.2 Rendering VEs: Constraints, Force Fields and Textures

We render a virtual environment in several steps, described in context of Figure 1.1's example: Harold's handheld car travels along a drawn road. When the car passes over the road edge, this interference is simply computed as a normal spring force.

First, we define the VE in the computer using our Python API. Then, in a 10ms loop (limited by transmission speed not computation): (1) surface-contact position is measured, amplified, filtered, and differentiated for velocity; and (2) sent through a forward kinematic model:

$$\begin{Bmatrix} V_x^S \\ V_y^S \end{Bmatrix} = R(\Theta)Y(\alpha)P(\beta)\begin{Bmatrix} V_x^d \\ V_y^d \end{Bmatrix} \tag{3.7}$$

.

(3) Interaction with the VE involves collision with VE elements, determining depth of interference, and computation of the resulting interference force.

(4) VE force is transformed into desired F_S and V_S by inverse kinematics (Eq 3.7); then through a voltage constant to command force on respective actuator pairs.

Braking and tactile display: The most effective braking occurs when the rollers both turn inward, driving the contact ball towards the trackball sensor. In this situation, the contact ball is actively jammed at four points (surface, trackball and two bearings). This locking effect can be leveraged for Cobot-like virtual constraints [40]. However, the same effect may occur inadvertently under excessive hand pressure on the device, which squishes the ball and prevents the drive from functioning. Users need to maintain a light normal force to achieve the intended force feedback.

When the rollers both turn outward, there is some braking but also vibration. More deliberately, rapidly alternating push/pull directions can produce up-down contact ball vibration, which can be modulated for texture rendering.

3.6 Design Review

Table 3.1 summarizes the degree to which this first prototype hereinafter referred to as "MagicPen" addresses our ballpoint drive design considerations.

We measured properties such as mechanical impedance and 2-axis forces (®BOSE ElectroForce TestBench). We defined velocity (V_{max}) as how fast the pen can move with motors saturated and no force resistance, and measured it by holding the ballpoint device stationary over a second, inverted trackball. We found that this mechanical and electrical hardware design can convincingly support the basic virtual environment primitives described in Section 3.5.2, and it can can keep up with brisk hand motion (20 mm/s).

However, there is ample opportunity (and means) for improvement. The most critical mechatronic improvements are drive train impedance, and inter-processor communications speed (to reach 1ms rather than current 10 ms).

Others will require straightforward engineering iteration: *e.g.*, eliminating data and power tethers and reducing drive system bulk while further reducing cost. Parts cost are already comparable to other DIY examples [149][67], custom electronics will drop overall parts cost by 20-40%. For more technical demonstrations and results, please see Appendix A.

Table 3.1: Design Considerations: Ballpoint Drive V1.0

Performance	F_{max}: 1N force (x_d,y_d). Can move the device body and convincingly drive the user's hand. V_{max}: 20mm/s across 2D surfaces, unrestricted. *Mechanical impedance*: 35 Ns/m (measured with 1 cm displacement at 5 Hz). This is greater than ideal haptic backdrivability.
Accessible Cost	Device parts cost are ~$101 USD: $50 (electronics), $12 (drive mechanism), $34 (orientation sensor), $5 (trackball sensor). Electronic components currently dominate cost, and can be reduced 40% through integration and more optimized choices. The drive mechanism would most benefit from component upgrades, as well as investigation of alternatives (*e.g.*, belt/pulley or metal gears to reduce noise).
Rendering Pipeline	*Rate:* We currently can render force commands at 1000Hz, and communicate them at 100Hz. *VE:* Our drive mechanism can rapidly switch operation mode for specific rendering conditions, such as a virtual wall, elastic behaviour and textures.
Ergonomics	Device weight (sans battery or Arduino circuitry) is 50 grams. Device is borderline adequate and will benefit from further miniaturization.

Chapter 4

PHysically Assisted SKetching

Preface – *In Chapter 3, we introduced MagicPen and its potential applications in learning. For our first application, we study the benefits of force feedback for assisted sketching. We investigate the opportunities in using different force feedback interactions and propose a framework that covers related research and unveils the existing gap in the design space. We realize that control sharing is vital to maintain creativity and authorship while collaborating with a computer in assisted sketching tasks. We use the pen-holder form factor that we introduced in Chapter 3, and instantiate our framework. We identify the conceptual activities and core interaction concepts that we need to support assisted sketching with our MagicPen. We validate our framework and our approach to implement it, by running a user study with both experts and novices* [1]*. We reflect on the results and the opportunities for future improvements. This chapter builds the foundation for the Design process which we will explain in Chapter 6. The material in this chapter is taken directly from sections: Introduction, Phasking framework, Evaluation, and Discussion of (Kianzad* et al. *2020* [2]*) while the System section is presented in Chapter 3 and the technical part of the Background section is in Chapter 2.*

[1] The approval for this study was obtained from the University of British Columbia Behavioural Research Ethics Board, approval ID (H14-01763).

[2] S. Kianzad, Y. Huang, R. Xiao, K. E. MacLean, "Phasking on Paper: Accessing a Continuum of PHysically Assisted SKetchING", Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, April 2020, pp 1–12.

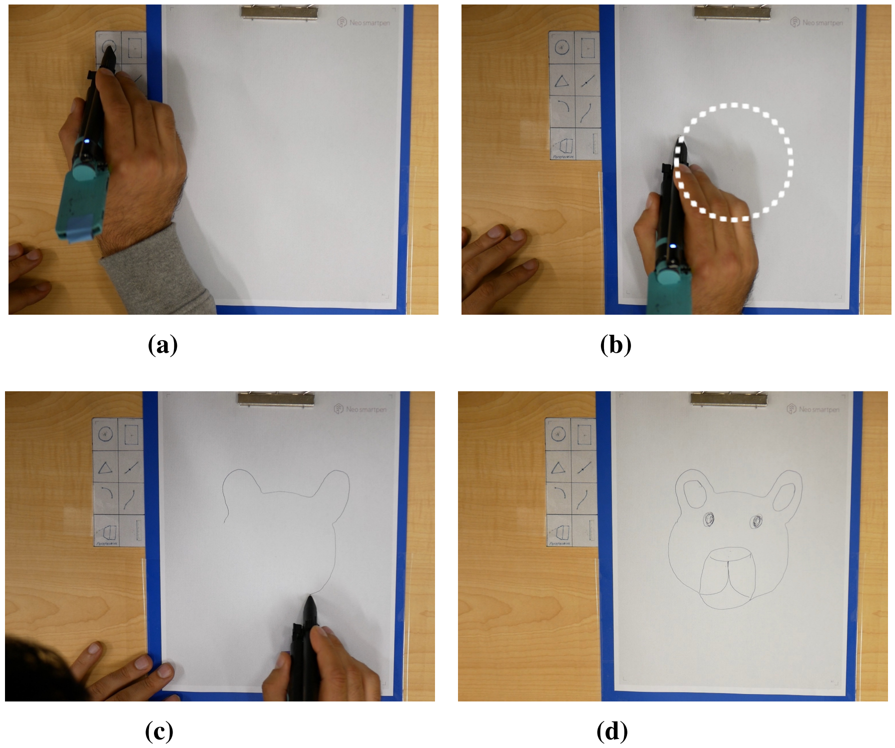

Figure 4.1: The steps of a phasking interaction. (a) The user selects the "circle" tool from the Phasking tool palette, (b) then selects the centre and one point on the circle to establish a circular *bring* constraint (dotted line). (c) The MagicPen actively *brings* the user along the circular path with its ball-drive motor, but the user can modify *control sharing* by applying pressure to the pen (Figure 4.9Right). This causes the system to scale down its constraint force, allowing the user to diverge from the path. (d) Phasking supports passive constraints as well as fully unconstrained drawing, enabling the user to quickly sketch out a cartoon character.

4.1 Overview

When sketching, we must choose between paper (expressive ease, ruler and eraser) and computational assistance (parametric support, a digital record). PHysically Assisted SKetching provides both, with a pen that displays force constraints with which the sketchers interact as they draw on paper. Phasking provides passive, "bound" constraints (like a ruler); or actively "brings" the sketcher along a commanded path (*e.g.*, a curve), which they can violate for creative variation. The sketcher modulates constraint strength (control sharing) by bearing down on the pen-tip. Phasking requires untethered, graded force-feedback, achieved by modifying a ballpoint drive that generates force through rolling surface contact. To understand phasking's viability, we implemented its interaction concepts, related them to sketching tasks and measured device performance. We assessed the experience of 10 sketchers, who could understand, use and delight in phasking, and who valued its control-sharing and digital twinning for productivity, creative control and learning to draw.

4.2 Introduction

From scribbles to detailed, elaborated productions, sketching is both intellectual play and can help us form, develop and communicate our thoughts, a key part of conceptualization. Pen-and-paper sketching is direct, improvisational, expressive, resists distraction, and may promote deeper cognitive processing [163]. The freedom and functional control afforded by physical drawing is unmatched in electronic media; but paper sketching lacks digital enhancements, and it is laborious to move fluidly between paper and digital media.

Meanwhile, freehand drawing is poorly supported in graphical and CAD (Computer Aided Design) environments, as evinced by many professionals' preference for paper. Capturing the subtlety and nuances of physical drawing and painting remains elusive [219].

Active force feedback is an intriguing approach to supporting on-paper drawing, with different opportunities for expert and novice sketchers. But active force feedback (rather than brakes, which cannot actively guide, or tactile vibrations, which cannot even constrain) usually entails a grounded device [201] with a fixed and impractically limited workspace, often costly. What if a user wants to roam with their physical drawing support, accessing guidance and constraints for everyday drawing on arbitrary media, or to support big strokes on big surfaces?

The goal of this chapter is explore the use of active force feedback in paper sketching, to provide a user with the digital support they need in the moment while maintaining their originality, authorship, and continuity of expression.

4.2.1 Physically Assisted SKetching with Variable Control

Phasking enables a user to access drawing supports in a continuum from full physical guidance to expressive freehand sketching, on their media of choice (Figure 4.2). It is free-roaming (untethered and portable); and our demonstration prototype is constructed of low-cost commodity components and DIY construction. It is based on two key interaction concepts.

Constraints: Phasking's bring/bound system captures and extends the range of assistance explored in previous work, by accessing an interactive, constructable virtual environment (VE). The user can place structures in the VE to constrain their movements both passively (*bound*) and actively (*bring)*, and can interact with constraints and the sketch itself.

Control sharing: Users can fluidly move between levels of assistance, simply by bearing down on the pen tip. Phasking differs from adaptive force feedback [18], where the system chooses the degree of assistance based on user performance. Phasking gives this choice to the user.

4.2.2 Usage Scenarios

A wide spectrum of users sketch, in many contexts. We consider situations where phasking will be valued, to prioritize feature development, and consider them in our evaluation.

Rapid technical sketching: Professional architects, engineers and other designers sketch copiously, rely on it for conceptualization and communication, and move between paper and digital media. Many experts value drawing assistance when on paper, whether to construct a perfect circle or perspective, as evinced by their heavy use of physical guides. Some find physical tools cumbersome and "in the way", *e.g.*, wanting to draw on both sides of a ruler. Large-scale drawing (*e.g.*, during public communication), is a situation where even skilled sketchers may have difficulty with alignment and clean curves. Finally, professionals spend a long time *learning* to draw. Phasking's large workspace, portability and digital-twinning capabilities could assist in all these contexts, both manually and by off-loading some cognitive demand.

[Re]Learning to sketch: Learners may be children, hobbyists or recovering stroke patients. All may need assistance in drawing a straight line, getting proportions right for an animal or a face, or figuring out perspective. There is evidence that CAD can help creative self-efficacy [203]. Learners' needs might range from manual control (*e.g.*, drawing a circle), to the cognitive challenge of using proportion. Learning is best accomplished on paper, with its friction, focus and room to spread out, but would benefit from computational supports. Phasking's constraints can assist learners and patients in reinforcing motor programs, and advanced skills like perspective drawing. Once learning is achieved, they might continue to phask in more expert ways, or no longer need it.

Artistic 2D Sculpting: Creative expression is inspired by constraints [135]. Physical constraints that can be stretched and violated stimulate a resonant collaboration between user and system [135, 207]. Phasking is founded on collaborative control-sharing. Artistic sketchers can follow a basic shape, or creatively modify it, by altering control authority. To draw a face instead of that circle, they press down to take control to form an ear or a sketchy line of hair. Or, they can pull and bounce off an active node to draw sweeping trapeze-like trajectories, bound to the node with an invisible elastic string.

4.2.3 Contributions

1. *A conceptual framework for phasking* which highlights fluid transition of control sharing from assisted to freehand drawing, making use of both bound and bring constraints.
2. *A force-feedback digital manipulative* capable of implementing this framework. We made major mechatronic, ergonomic and control extensions to a previous ballpoint drive display to enable screen-free, shared control phasking.

4.3 Background

Our work is founded in manual and computer-aided drawing practices, virtual and augmented environment creation and manipulation, haptic force feedback and that field's knowledge of control-sharing and past frameworks for sketching support.

Figure 4.2 arrays the mediums and examples mentioned here on the spectrum

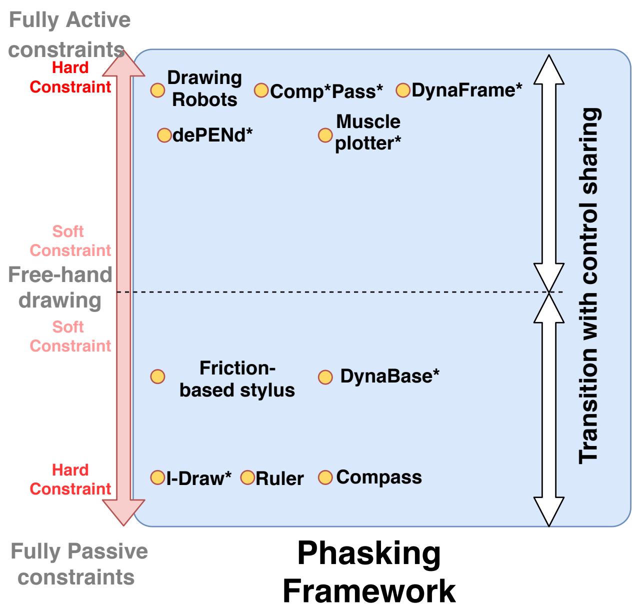

Figure 4.2: *A one-dimensional model of sketching control authority.* The *y*-axis denotes strength of a given system's constraint, *active* up / *passive* down. The *x*-axis has no meaning. Existing types of assistance are located approximately in this space, both generic (*e.g.*, a physical ruler) or published works (denoted with *, see Related Work). Phasking can fluidly access all points on this axis, by bearing down on the pen-tip while drawing.

between fully freehand and fully computer-aided drawing. Vertically, this figure highlights how *control sharing* is related to type of drawing. Most examples occupy only one point in this space, whereas physically assisted sketching, as presented here, can theoretically cover all of it.

4.3.1 Non-Haptic Assistance of Digital and Manual Drawing

Professional CAD tools are increasingly accessible, online and learnable: basic ideas of parametric drawing have high penetration for even minimal expertise. But despite many conveniences (pen type/color, copy/paste/undo) the experience is still fundamentally different than pen-and-paper drawing due to tactility and limited canvas.

Graphical drawing systems can provide visual corrective feedback for drawing or give stroke suggestions, *e.g.*, iCanDraw to sketch a human face from a source image [50], and ShadowDraw for high-level arbitrary objects [133]. These systems typically restrict users to graphical screens.

Paper-oriented digital styli, such as Anoto and Neo smartpens, feature real-time digital capture of handwriting and translation to digital task, requiring use of watermarked paper (*e.g.*, pre-printed with microdots) [97, 98]. While capturing natural drawing, this approach cannot offer added support *during* sketching. Our conceptual prototype incorporates a Neo Smartpen with its pen-tip vision system and watermarked paper as a convenient way to mockup position localization.

For *visual* guidance, PenLight [208] combines an Anoto pen with a miniature projector to add information to pen and paper interaction, but faces a technical barrier of image stability In virtual reality, Nomoto *et al*. present a "corrected" sketch which encourages the users to configure their own hands appropriately for drawing a shape [167]. These are promising approaches but do not offer physical constraints. In phasking, real forces on the hand – in the real world, not a VR headset – convey drawing suggestions.

4.3.2 Haptic Support of Pen-and-Paper Sketching

Passive Constraints: A ballpoint stylus able to impose passive constraints by *constricting* a rolling contact ball by means of electromagnetic or mechanical brakes has some similarity to our ballpoint drive and was used to render roughness on 3D objects, but cannot provide active forces [43]. Comp*Pass [166] offers a semi-active solution using DC motors; however, a user is not actively involved in sketching. With I-Draw, a cobot-type drawing assist for passive constraints [59], the authors "explore the seamless switching between guided and freehand modes," as do we.
Active Guidance: Muscle-Plotter generates force feedback to the hand by electrically stimulating the user's own muscles [142] While creatively satisfying the criteria of free-roaming, it has drawbacks of 1.5 DOF, and a potential for temporal adaptation by muscles [110][240].

dePENd [231] exploits the ferromagnetic property of pen ball-tips to assist sketching by providing directional force feedback on regular pen and paper interaction. The main drawback of magnet-based haptic assisted sketching devices [140] is the tradeoff between backdriveability and perceptible force levels. Increasing magnetic coupling provides higher forces, but draws the pen tightly to the interaction surface (higher normal force) and makes it difficult to move freely.

I-Draw, dePENd, and Muscle-Plotter each demonstrate a concept of guided drawing (using actuation to turn a user's hand into a computer-guided drawing implement in a screen-free context), one passive and the other active. They are a departure point for the contributions described here.

4.3.3 Using Haptics to Facilitate User-System Control Sharing

Control-sharing with haptic systems has been utilized for physical therapy [136] and handwriting control [215]. In driving, haptic shared control can improve speed and accuracy of human / system collaborative tasks, and lower the need for visual involvement in control effort [3, 76] It has also been used to support expressive drawing on a screen, *e.g.*, Snibbe's Dynasculpt and GridDraw [207]. While I-Draw [59] is framed in moving between freehand and supported drawing, its passive nature does not provide an ideal mechanism for doing so.

We have drawn from these functional and expressive approaches to form our own control-sharing, which prioritizes simplicity and intuitiveness in modulating control authority in instances where users need a collaboration rather than a binary choice. The MagicPen's capabilities support this.

4.3.4 Frameworks

Steimle *et al*.'s framework of non-sketching pen-and-paper interactions separates conceptual activities (annotating, linking, tagging) from core interactions (inking, clicking, combining, associating) [211]. While its domain differs, we are inspired by its approach in our own support framework.

I-Draw presents an initial framework for passive guidance, of interaction primitives allocated between physical (guided and freehand drawing) and digital (digital manipulation) spaces [59]. We re-organize and extend it with the capacities afforded by active haptic guidance.

4.4 Phasking Framework

The added capabilities of a fully force-performant but free-roam, screen-free device has several implications. First, the availability of active, omnidirectional forces in a handheld format permit fundamental changes in interaction, notably active guidance and control sharing. This physical support can work in both graphical and screen-free contexts, widening scope and altering how digital-physical transitions can occur. Finally, the active force's scalability means that constraints and guidance can be modulated, from hard to soft. Together, these necessitate a deep revision and extension beyond past conceptual framings (*e.g.*, [59, 211]).

Like [211], our framework articulates *conceptual activities* that users need to do, elaborated in (I) below; leading to *core interactions* that support them (II): bound and bring-type constraints and variable control authority (constraint hardness).

Figure 4.2 shows how bound/bring constraints interact with shared control in the phasking framework.

4.4.1 I. Conceptual Activities

We articulate the foundational activities which our framework needed to support, as elements that mediate a dialogue between user and system in Table 4.1, a potentially extendable list. These emerged from our observation of users' expectations formed through interacting with conventional tools, as well as consideration of the basic operations of freehand paper drawing, CAD and virtual environments.

4.4.2 II. Core Interaction Concepts

These concepts demonstrate how phasking's key conceptual activities are supported. We use a constraint-based virtual environment (VE) which a user constructs then sketches within, with tools created by drawing a palette on the paper's margin.

(a) Constraints – Bounding and Bringing

Constraints can be expressed as a gain on an error function (of position, velocity or other parameters): $u = K(x_{des} - x_{act})$. With active force assistance, phasking constraints can passively *bound*, or actively *bring* (Figure 4.3). They do both to varying degrees (Control Sharing concept, below).

Bound: Movement is free up to the boundary, then constrained. A binary boundary

Table 4.1: Core conceptual activities of the phasking framework.

Activity	Description
Free-draw marks	Create marks manually and at will, optionally within user-set and system-maintained boundaries.
Create objects	Form basic shapes on command (*e.g.*, parametrically specified). Produced objects may conform perfectly to the digital guidance, or the user can overcome guidance to construct personalized or expressive variations.
Place & arrange elements	Receive assistance as to where, how large and what angle; *e.g.*, perspective drawing, or sizing different regions of a multi-part sketch
Interact with active constraints	Set up constraints (*e.g.*, attractive nodes, or lines and curves to push/pull against) for modulated creative control in variably-guided drawing.

(no shared control) could be implemented with a passive force device, *e.g.*, a brake, because it just prevents the user from going somewhere. Examples of *bound* constraints include one-sided walls, and path constraints which the user can traverse at will: the constraint blocks path departure, but allows free movement along it.
Bring: A force field draws the user in a particular direction or rate, and requires an active force feedback device; it always entails active guidance[3]. Examples of *bring*

[3] *Bring* contrasts with what is called guidance (but is passive) in some related work (*e.g.*, I-Draw),

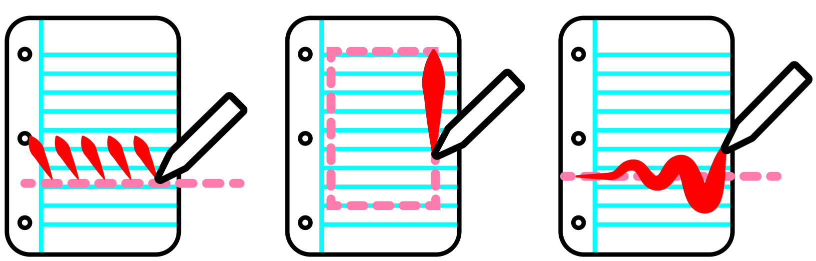

Figure 4.3: The framework's (a) *bound* and (b) *bring* constraints, and (c) the concept of *control sharing*, where the user can diverge from a guiding line by bearing down on the pen.

constraints include point magnetic attraction or repulsion (snap); and spatiotemporal and temporal trajectories, in which the user is guided to traverse a path in time and space respectively.

As with a ruler, passive VE elements are only felt when the drawing tool touches them. Active elements, *bound* or *bring*, can be felt at a distance, as a force field.

(b) Control Sharing

In phasking, what's shared is control authority ("who gets to drive"), under continuous user control (*e.g.*, the pressure with which the drawing tool is squeezed or pushed into the drawing surface). The constraint can vary between absolute and soft – a suggestion or a jumping-off place, *e.g.*, if one wishes to draw a wiggly line along a path (Figure 4.3c). This scaling is available for both *bound* and *bring* constraints.

Control sharing can be implemented simply by changing the control gain on the error between pen and constraint. A wall can be softened, and an actively guided geometry (such as an oval) can be sketched into something more expressive and detailed, like a face. More complex implementations are available to address system stability issues [3, 63].

(c) Tool Selection

Like other digitally assisted drawing systems, phasking is modal. We deliver the function of tool selection with a paper tool palette with hand-drawn (and extemporaneously creatable) icons that the pen's vision system can recognize (see *System*). Because interactions are brief, tool selection also supplies modal awareness, together with the device's physical response.

(d) Constructing Constraint Environments

In free-hand drawing, a user sets a passive boundary with a tool, like placing a physical ruler to help draw a straight line or set a CAD drawing plane.

Phasking constructs and tracks constraints through a virtual environment: the user draws the environment within which they then operate. For example, the user places a virtual bound (*e.g.*, a line, circle or channel) or an attraction point for a *bound* or *bring* constraint, respectively. The VE is portrayed to the user through both the visible marks on the drawing medium, and what they feel.

---

in which a passive mechanism such as a a Cobot [40] or brake restricts movement in some direction.

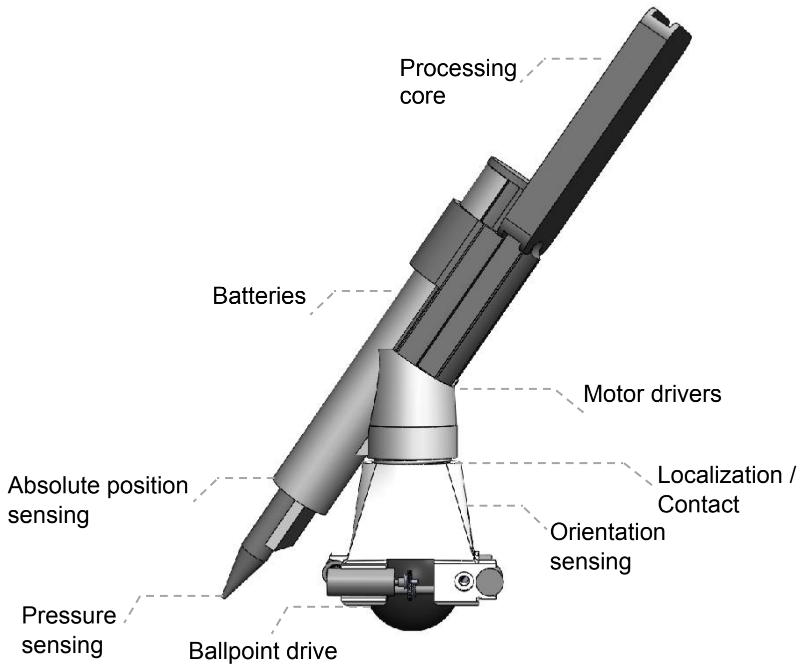

Figure 4.4: The MagicPen: an untethered ballpoint drive produces forces by driving a contact ball over a surface. Friction between the ball and surface prevent slip, and provide a "ground" back to the user's hand.

Constraints support expressive drawing, and force feedback has been used for this [207]. But when working screen-free, constraint creation *is required* to access assistance. Here, our VE is a basic functional implementation, but the approach opens other design spaces as well.

and paper-based operation. Here, we overview the full system for the reader's benefit, but focus on novel or modified elements.

4.4.3 MagicPen Mechatronics

We implemented the MagicPen's mechatronics, system architecture and controls, and phasking primitives to assess the concept's feasibility and usefulness. The device reported here significantly extends a previous basic demonstration of the ballpoint drive mechanism as a 2D force-feedback display (Figure 3.1a), as required for this application. The detailed view of the system is presented in Figure 4.4.

Mechanism: Untethered 2D Forces Via Ballpoint Drive

In the ballpoint drive [113], pairs of opposing motors drive a surface-contact ball to create directional force-feedback, generated between ball and arbitrary two-dimensional surface (Figure 3.2). We completely re-engineered the ballpoint drive mechanism to improve backdrivability (freedom of unpowered motion), reducing passive impedance by 51%. We also reworked the gear drive to achieve higher force without slip, significantly reducing vibration and skidding. We customized the drive train with low-cost commodity micro gears, motors and a surface contact ball. To achieve required backdriveability and power transfer, we iterated component configuration, size and material properties (Figure 4.5).

Gear drive design : Ensuring non-slip coupling between motors and surface contact ball can add impedance to the drive train, which then degrades rolling free-ness when motors are not actuated. We found a solution by matching contact-ball material properties with cog size. A 1-inch diameter rubber ball with tensile strength of 144 MPa, with metal gears with diametral pitch of 187 teeth/inch [53] gave the best results. The contact ball has been sandwiched between the metal gears. The metal gears' teeth penetrate the contact ball's rubber body just the right amount giving the contact ball enough room for rotation while creating non-slip contacts with the gears.

Drive motor coordination: We can achieve optimal control over the surface contact ball with four motors (as opposed to two drivers and two passive castors; or a 3-point contact which cannot reach the same control space). To generate a rolling movement, each opposing motor pair works together to apply a balanced torque to the ball. While it is possible to electrically connect the opposing pairs, we found out that for fast movements when the motors are not in active mode, paired motors work as a generator and produce back-emf voltages which pass through the opposing pair. To avoid the consequent resistance, we drove each motor separately.

Computing processor, motor drive and communications: The primary processing unit (Raspberry-Pi Zero W, or RPi) takes sensor data, updates a state model, then computes proportional-derivative (PD) control commands and sends them to two Pololu DRV8835 dual-motor driver carriers (one per axis).

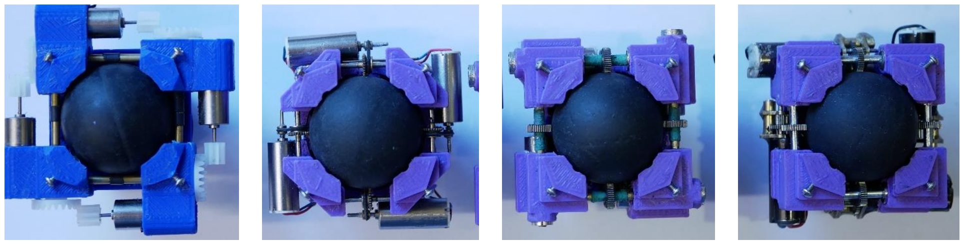

Figure 4.5: MagicPen ballpoint drive design iterations. From left (earliest): (a) Early version with plastic gear between motors and rollers, and rubber connecting ball between roller and surface contact ball. (b) Pulley belt connection between motors and rollers, and clock gears with micro cogs between roller and the contact ball. (c) Metal gears between motors and roller, for a lower gear ratio of 1:1. (d) Similar mechanism with a customized higher gear ratio of 4:1.

We sample micro-trackball velocity, integrated to get position at 5 kHz (RPi external interrupts), and BNO055 for orientation sensing at 100 Hz. The Neo Smartpen sends its data (x,y position and pen-tip pressure) to the RPi controller via Bluetooth Low-Energy at 100Hz. We built a custom Linux driver for the Neo's BLE protocol to reduce latency and enable custom features (e.g. on-demand beeping).

State model and closed-loop control: A 1kHz control loop checks for a command, samples internal position and orientation, receives x,y and pressure data from the Neo (100 Hz), and optionally sends it to the monitor. Control then branches:

Table 4.2: Control steps.

Free mode	No command is running; wait for next iteration.
Command mode	If new command registered, collect command parameters parameters (position taps). Then: ▹ Update command target reference. ▹ Adjust absolute contact position estimate based on pen orientation, with internal position change. ▹ Compute motor command command using a PD controller on the error signal. ▹ Output motor command to PWM motor.

4.4.4 Implemented Drawing Support Features

2D CAD software implementation: Our new untethered design required a custom, lightweight CAD platform for the RPi, and custom low-latency hardware drivers for the ballpoint drive and digital pen.

To implement the MagicPen's drawing functions, we developed a simple 2D CAD software system using Python and the PyQt4 library (a Python interface for Qt, a popular cross-platform graphics library) [139], which runs on the pen's RPi controller. With this software, the user can either free-draw, or access primitive CAD functions by clicking on a paper tool palette (described below).

The CAD software also implements a graphical user interface view (referred to as the *GUI monitor*), which can be displayed on a screen connected by HDMI cable to the onboard RPi for debugging. The GUI monitor view shows real-time updates of the user's drawing, and provides additional functions such as saving the digitized drawing locally and changing the color and thickness of the pen.

(a) Bound/Bring Constraints for 2D Geometry Construction: We implemented eight phasking primitives (Table 4.3): seven to construct basic shapes or constraints, and a perspective function for use with other primitives.

Each uses *bound/bring* constraints, with force guidance modulable through control sharing (below). Jointly, these primitives implement all of the core conceptual activities of Table 4.1.

(b) Sharing Control Authority: When drawing, the user can start with one of these primitives and deviate from the pen's guide by applying a small force to sketch more complex shapes (Figure 4.10).

(c) Tool Switching: paper-based tool palette: Phasking requires extendable access to the drawing primitives. We used a paper palette as a simple physical access point; the user can draw tools on the sheet, selects a command by tapping on a box, then taps on the drawing to define parameters (Table 4.3). Figure 4.6 shows the watermark-paper implementation. The tool palette has an added advantage of logging user commands, as one route to saving geometry; *e.g.*, for copy/paste on paper, or a screen-based reconstruction.

Table 4.3: Phasking primitive descriptions. Each operation begins by touching the corresponding icon on the tool palette.

Primitive	Description	Illustration
Line, Ruler	(1) Touch end point (2) Touch starting point (3) *Ruler*: draw along invisible barrier between the two points *triangleright Line*: MagicPen *brings* to end point	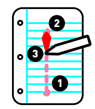
Triangle, Arc	(1-3) Touch three points to define Triangle or Arc ▹ MagicPen *brings* across triangle edges, or arc.	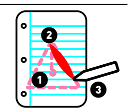
Circle, Rectangle	(1) Touch center (2-3) Touch radius (circle) or top left corner (rectangle) ▹ MagicPen *brings* to endpoint.	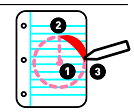
Bezier spline	(1-4) Touch at least four control points ▹ A cubic (4 points) Bezier curve is defined (5) Touch curve ▹ MagicPen commands motor velocity according to the tangent line to the curve at each point.	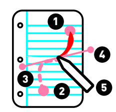
Perspective function	(1) Touch vanishing point (2-n) define any geometry, *e.g.*, *Rectangle* (center, corner) ▹ MagicPen draws the object (*e.g.*, *Rectangle*) in perspective	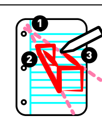

4.5 Evaluation

Our evaluation objectives focused on conceptual viability: we needed to know whether the ballpoint drive approach could perform well enough to support phasking operations, and to get insight into users' experience of phasking.

We tested force-feedback standards of force and position control and disturbance rejection. These address whether the novel drive can localize itself with its dual sensing system even as the user's hand rotates the pen's axis (pitch and yaw), follow a commanded path, provide a usefully large commanded impedance, and reject disturbances, all fast and smoothly enough for at least moderate-paced drawing. We sought performance that would let us try phasking out, an assessment to be rendered in part by our user study. These tests also set a benchmark for comparison with future progress.

4.5.1 Performance Characterization

Prior to involving participants, we evaluated the MagicPen's mechanical and control performance.

Test 1 – Force generation: We measured passive and active 2D forces with a BOSE ElectroForce TestBench®, which held the MagicPen with two arms (Figure 4.7a).

Passive force, step response: Low mechanical impedance is crucial for control sharing interactions. We recorded MagicPen's resistive force while unpowered, while one BOSE arm imposed an 8mm sine-wave position displacement to the MagicPen at 0.5 Hz for 5 periods. We found an impedance of 37.5 Ns/m, computed by dividing the recorded value of mechanical resistance force over the speed. This is approximately half that measured for the previously reported version of the ballpoint drive (77 Ns/m) [113].

Active force, step response: We measured the force the ballpoint drive produced in response to pulsed drive input (~ 0.2Hz), while held isometrically between two measuring load cells, one on each arm (Figure 4.7b). The pen generated up to 1N in continuous force (close to peak output) without slip between contact ball and gears, after which the drive gears slipped and the ball started spinning.

Active force, sine response: As seen in Figure 4.7c, generated force closely followed a continuous sine pulsewidth modulated (PWM) voltage input. There was minor

Function	Palette	Function
line (*bring*)	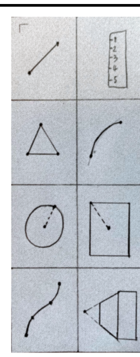	ruler
triangle		arc
circle		rectangle
B-spline		perspective

Figure 4.6: Paper tool palette, which can be hand-drawn and customized, or printed on a full sheet or slip of paper.

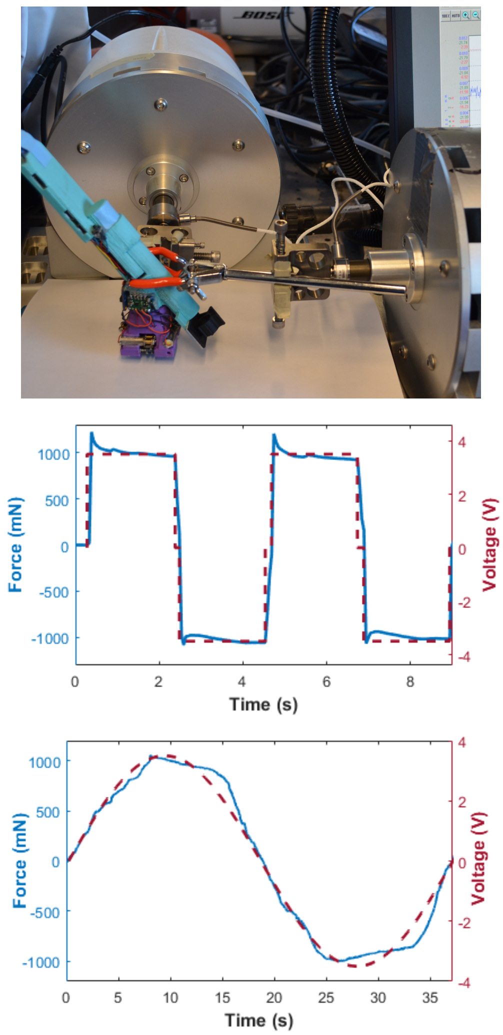

Figure 4.7: Force generation performance. (a) BOSE test setup. (b) Max output force response to PWM voltage pulses. (c) Force-tracking response to a slow sinusoid of PWM excitation.

nonlinearity at higher voltages, as excitation approached motor saturation.

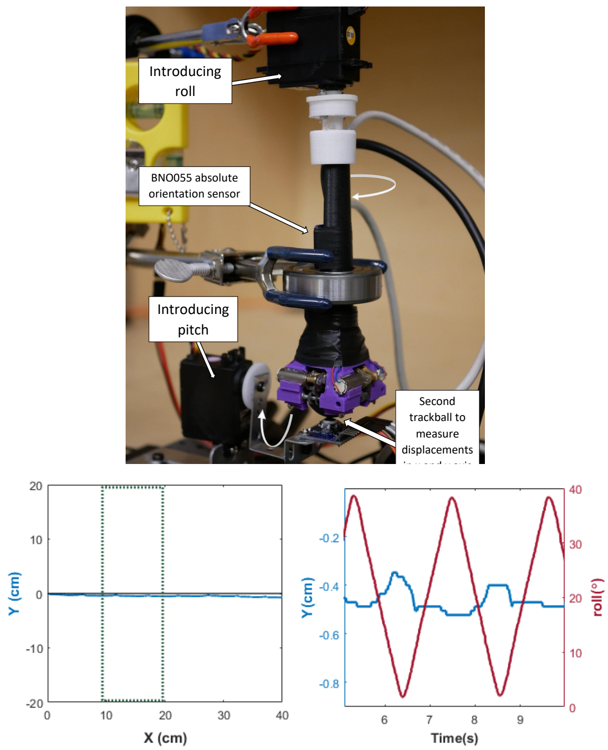

Figure 4.8: Position and disturbance tracking. (a) Test setup. The MagicPen's shaft is held in a vice, with roll and pitch disturbances applied at the top of the shaft and through the contact ball, respectively. The onboard processor controls the contact ball trajectory using orientation (roll/pitch/yaw) and internal ball motion. An external trackball measures ball movement relative to a global reference. To assess performance, we compare error between command and externally measured trajectories. (b) Test 3 (Disturbance Rejection) results shows the system's response to rapid yaw (0.5Hz, 35º peak-to-peak over 22 seconds), mimicking significant wrist rotation.

Test 2 – Position control: We required the MagicPen to follow a sine wave trajectory using only the internal micro-trackball (relative position sensing), measured with the setup of Figure 4.8. This test demonstrates the ballpoint drive's capacity to achieve agile, omnidirectional control in the absence of disturbances. We chose a sine wave target to capture a full range of movement in a 2D plane, and used a steady roll offset of 10°. This means that for the surface contact ball to, *e.g.*, move along an x-trajectory, it would need to adjust x,y motor commands using orientation data rather than simply turn the x-motor on and the y-motor off. Here, the task is to follow an x-y sinusoid. Due to the low impedance of the second trackball, the contact ball rolls at full speed (20mm/second) which slightly reduces the accuracy of the controller.

The MagicPen followed a 4.0cm peak-to-peak sine trajectory with 20% initial error and 40.0cm path length, using only the trackball (relative position sensing) corrected by orientation data for position control, with an error (mean squared distance to reference) of 6.78 mm (std 4.88mm), or 1.4%.

Test 3 - Disturbance Rejection: The controller needs to compensate when the stylus is twisted due to movements of the user's hand, adhering to a straight x trajectory with no y deviation. To do this, it just continually change motor velocity to maintain contact ball motion in the x direction alone.

Figure 4.8b shows the system following a line as we sweep through roll angles of 3-38°(a range observed for handheld usage) with mean square error 6.89 mm (std 2.29mm) 1.3%.

Introducing a pitch angle of 0-25° can reduce the speed in the x direction up to 10%, as the surface contact ball rolls on a smaller circle. This can be compensated either by multiplying a $cos(\theta)$ coefficient [113], or using absolute position sensing.

Test 4 - Sensor fusion: By using the absolute position sensing from the NeoPen, the controller can compensate for the offset and achieve higher precision. Figure 4.9 shows (left) the performance of the device in drawing a circle with a hard *bring* constraint, while a user gently holds the MagicPen; and (right) with control sharing activated while the user authors creative modifications to the basis circle.

4.5.2 User Evaluation

We conducted two user studies: (a) A performance assessment with novices (N=7) to assess usability and impact of core phasking interactions on normal users. We measured users' deviation from a predefined trajectory (error) and investigated their pressure profile to better understand the shared control (SC) concept in practice. (b) An interview evaluation with domain experts (N=3, architects who sketch in their

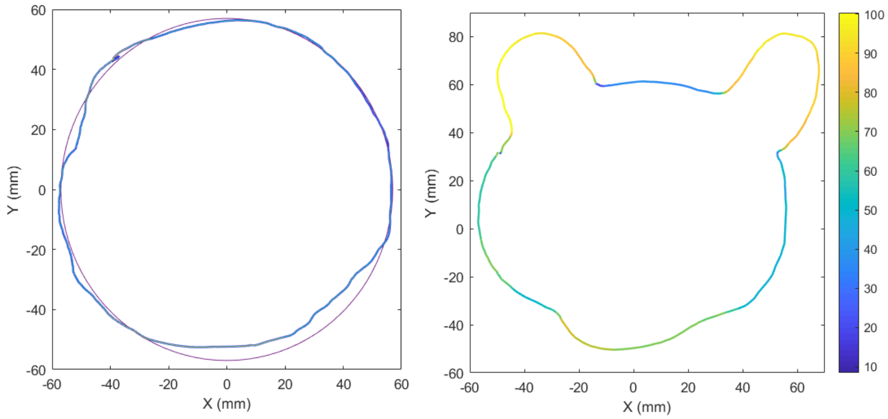

Figure 4.9: Human's aided circle following. (Left) System control: with control-sharing off, we fuse absolute and local position sensing to guide user P1 in drawing a circle. The deviation is a result of the pen rotating in P1's hand. (Right) With control-sharing on, a user violates a *bring* circle guide to draw a bear's face by pressing down on the pen. The color scale indicates applied pressure (yellow at 100% user control, blue when user has relaxed and is letting the system drive).

professional work) to assess value, fit to needs and potential for a enriching user experience, based on a functional conceptual prototype.

We identified domain experts as those who had professional training in hand sketching and usually use hand drawings in their day-to-day work. There was only one exception among our novice drawers, who took a preliminary drawing course around six years ago, but he has not been practicing drawing since then.

Procedure, both studies: We collected profile information on how participants used sketching in their work, and their personal attitudes to it. After a system familiarization session, we asked participants to perform several tasks (Table 4.4) with the Neo Smartpen on its own, and with the MagicPen. Neo/MagicPen order was counterbalanced by task and participant. Tasks were performed in the same order for each participant, as a progression in complexity (please check Appendix B for the user evaluation's results).

After the tasks, all participants were asked to complete a Likert scale (1:7) on: (1) *How likely are you to use this system again in the future?*; and, (2) *Would you*

Table 4.4: Evaluation tasks, by execution and complexity.

Concept	Task
***Bring* constraints**	[1] Draw a straight line [2] Draw a rectangle [3] Draw a rectangle in perspective (novices) [4] Draw a circle (novices)
***Bound* constraints**	[5] Draw diagonal lines meeting an invisible barrier (coarse cross-hatch on a line)
Shared control	[6] Draw a sine wave as MagicPen guides along a straight line (pulling towards the guide) [7] Draw a sine wave as MagicPen sets an invisible line barrier at the center of the sine wave (resisting their crossing of it).

use feature again (assume a more refined version of the tool)? for each of Line (Bring), Line (Ruler), Rectangle, and Shared Control. (3) *The movement speed and force is appropriate for me*.

This general procedure was the entirety of the Novice evaluation (30 minutes/session). Seven participants (aged 21-30, 4 female) had backgrounds of Computer science and Forestry.

Experts: Procedure for Qualitative Interview: Following the task-based interaction period, we carried out a 30-minute semi-structured interview with our domain experts, covering task experience, relevance to their professional work, and potential impact. A complete session for experts took 60 minutes.

Results

Quantitative Likert Responses (N=10): In total, participants performed 204 trials. Figure 4.10 shows the responses of all participants (novices and experts) to the survey questions. The Line function (ruler, or *block)* was the most popular feature, followed by SC, with *Bring* close behind and a strong interest in adoption. Rectangle was the only feature to receive any responses below neutral.

Table 4.5 shows task precision of manual drawing (NeoPen unmounted and used as a normal pen.) *vs.* phasking. Straight lines cover the basics; circle shows the system‘s ability to generate force feedback in a full range of 0-360 degrees, and rectangle+perspective a more complex task and guidance. Our data show shared

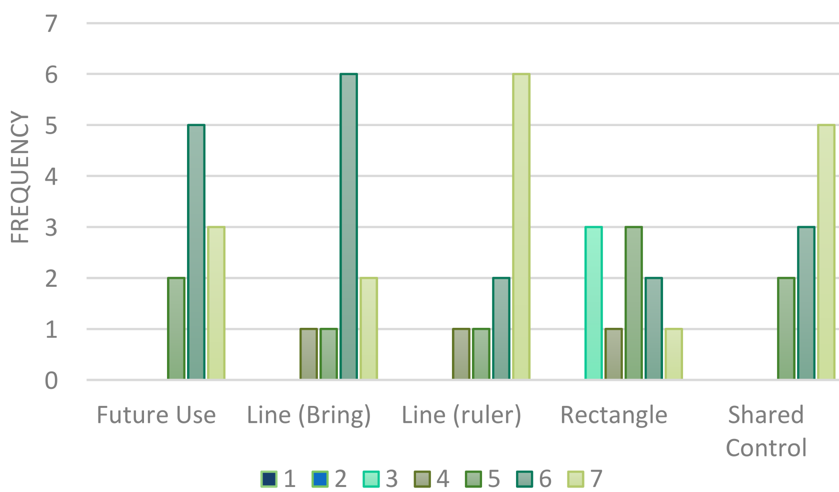

Figure 4.10: Users' Likert scale responses for the MagicPen(N=10, 7 novices and 3 experts; 7 is positive).

control reducing error, with both *bound* and *bring* constraints.

Experts' Profiles and practice: The three participants (2 female, all right-handed) had practiced in the area of landscape architecture for periods ranging from 3-12 years; E1 in outdoor, E2 urban and E3 residential design. None had prior haptics experience. All confirmed that their process started with hand-drawn conceptual sketches, an important aspect of both ideation and client and collaborator commu-

Table 4.5: Precision of manual drawing *vs.* phasking.

Task	System	Mean / std abs error, mm
Line	Neo	1.56 / 0.43, N=10
Line *(Bring)*	SC	1.38 / 0.33, N=10
Line *(Bound)*	SC	1.20 / 0.30, N=10
Circle	Neo	5.39 / 5.34, N=7
Circle *(Bring)*	SC	5.20 / 5.65, N=7
Rect	Neo	5.91 / 3.87, N=7
Rect *(Bring)*	SC	3.38 / 3.21, N=7
Rect+persp	Neo	9.21 / 6.86, N=7
Rect+persp *(Bring)*	SC	2.22 / 3.21, N=7

Shared control (SC) or just Neo Pen (manual: single point contact, no guidance).

nication, including in public meetings (E2). While estimating hand drawing as 10-20% of their entire process, after which continuity and precision required CAD, they wished for more:

> *My personal preference is hand drawing and sketching but I also like and appreciate the precision and the tools that other CAD basically provides you.* [E1]
>
> *It's not about preference. It's about the tools that we currently have.* [E2]

E3 additionally mentioned the importance of color in hand-drawn work, for impact in public communication.

Basic task performance, quality and experience: Participants rated appropriateness of speed and force at [5.5,7,5] on a 7-point Likert scale; E1 requested more flexibility in speed and force control. They could feel and understand the feedback, and control it to degrees estimated at 60-100%; E2 (60%) described "*an invisible barrier that sort of holds you ... keeps your line nice and tidy*". Some noted initial awkwardness, with greater comfort by the end of the tasks, and that some CAD functions had unfamiliar steps (*e.g.*, sequence of marking).

They found its use generally intuitive; E1 liked the ability to construct a perfectly straight line, and E3 noted that "*Line [ruler] part was really interesting because it really helped me to draw a straight line and it was the most interesting ... aspect of using it.*" [E3]

All were enthusiastic of SC's value and intuitiveness: "*an amazing sort of transition between a hundred percent computer drafting and hand drawing*" [E1]; "*in a way the device starts reacting smartly to what I'm intending to do.* " [E2]. E3 sometimes pressed too hard, then found SC less controllable. Precision (E1) and bulk (all) were identified as primary issues.

Relevance to professional work: All valued the potential to cycle between paper and digital work, in contrast to their present one-way transition. Of features evaluated (geometry, barriers, SC), all identified SC as most useful. Among specific widgets, they preferred Line (Likert responses [6.5,6.5,6]), with more mixed but still generally positive reception for Rectangle [5.5,3,6] and Ruler [4,5,7]. For screen-free and large-surface potential, E1 noted the difficulty of maintaining control on large surfaces (where MagicPen could help); E2 mentioned value in a public engagement process, and communicating extemporaneously with an audience. E3 wanted a phasking ruler for section and building elevations, now done by hand.

Overall impact and interest in adoption: Experts responded to the adoption question with [6,7,6]. Presuming a slimmed-down and more precise device, participants were positive on productivity (*e.g.*, by integrating paper sketching later in process). E3 indicated great interest in precise technical sketching, rather than conceptual work where roughness was fine, and liked the efficiency – "*you put away the ruler and then you have two things in one (pen and ruler)*". E1 predicted value in education, noting 10 years of training with constant practice.

4.6 Discussion

We have presented and implemented *phasking*, a form of computer-assisted drawing that brings a virtual environment constraint system to pen and paper, and allowed users to access a continuum of assistance (type and hardness of constraints) via fluid sharing of control authority.

We created this framework out of the varied purposes that people bring to sketching. Phasking requires active force feedback, because it entails active and passive constraints, and user-controlled gradation in, *e.g.*, a restoring force upon violating a helpful constraint. Paper sketching requires an implement that can operate screen-free. For large, free movements, the device needs to be free-roaming (untethered).

In creating the Phasking pen, a major extension of a previous ballpoint drive device, we focused on strength and backdrivability, attributes difficult to jointly optimize but crucial for feature rendering and unimpeded movement. Objective performance metrics provide a benchmark for future improvements. We implemented an essential set of framework primitives and a paper-based tool palette to access them, with which users can carry out a complete drawing task on paper.

We shared phasking with novice and professional sketchers. They could feel and understand the forces, and found strength and speed adequate while wanting more precision. They also told us of a strong desire to be able to hand-sketch more, which requires integration throughout their process, not just at the start. Because of the volume of their use, they valued physical supports for productivity and prized fluid control sharing; and suggested that geometry construction would have been valuable when they were learning to sketch.

Limitations: Rotation in the drawing direction (pitch) is constrained by two-point contact, addressable by delivering ink via the ball itself. Our evaluation revealed notable individual differences, particularly in magnitude and smoothness of force profile, signalling a need for training. While new users learn to deploy pressure to optimize SC use, graphical or auditory feedback will be valuable. A screen as training-wheels could also assist with learning CAD features. Finally, a nonlinear relationship between pressure and control share might work better, a topic of future work.

4.7 Conclusions and Next Steps

Despite some usability issues, which are an inherent part of early prototypes, we were able to assess phasking's potential and collect feedback points to key improvements.

With a full interactive experience in place, we have proved possible many other interesting functions within this framework, including modifying elements (*e.g.*, resize, rotate, amend) as well as copy, paste, undo; identifying free-drawn marks as parametric objects; combining objects into a virtual construct, and even simulating dynamic virtual systems. It is a small straightforward step to full *digital twinning*: modify it onscreen then bring it back to paper with guided tracing.

Phasking is too different from other digital tools to know its full potential. Being untethered, portable and self-contained, MagicPen can, with attainable modifications, be used on arbitrary surfaces. This could lead to a new way of 'drawing on a napkin', support blind mobility by revealing maps on a corridor wall, and allow drawing and playing with simulations on a whiteboard – an 'object to think with' [174, 190].

Chapter 5

Benefits of Force Feedback for Collaborative Grounding

Preface – *Whether touch sensory feedback can enhance a learning outcome is difficult to ascertain; long-studied, there is no conclusive result to date. Many researchers use theoretical and empirical pieces of evidences to support their research hypothesis. From the theoretical standpoint, embodied cognition and touch as an additional sensory input are the most popular theories to justify the benefits of physicality and haptics in learning. On the other hand, the empirical research was often devised to seek significant differences between the learning gain of a controlled group without haptics vs. a group with haptic feedback. In this chapter we discuss the findings and obstacles identified in previous work and offer a new lens to study the benefits of haptics in collaborative learning by focusing on the strategies that learners take to achieve mutual understanding and reflect on the learner's haptic experience model*[1]*. We used the stylus form factor (see Figure 3.1(a)) in this study. The chapter's findings have potential benefits in two fields of research, namely educational haptics and collaborative learning. Further, we use this chapter to ground the Explore process in Chapter 6.*

[1]The approval for this study was obtained from the University of British Columbia Behavioural Research Ethics Board, approval ID (H14-01763).

5.1 Overview

We investigate how the sense of touch can support the development of mutual understanding (grounding) between collaborating learners via *dialectical discourse*: a conversation in which differing points of view move to consensus through reason, here supported by evidence collected with a shared force-feedback environment. Using collaborative learning as a lens, we invited dyads to physically explore virtualized phenomenon related to physics concepts, and examined their strategies for using (a) force cues and (b) mutual understanding in learning tasks. We compared their behavior across three diverse low-cost haptic platforms, each paired with a learning environment and activity that exploited its best attributes.

In Study 1 (n=8), we confirmed the occurrence of haptically mediated grounding and identified patterns between grounding acts, haptic gestures and learner intentions. In Study 2 (n=24) we assessed collaboration dimensions through an objective lens focused on *haptic critical instances* (hCIs). Qualitative and quantitative evidence from these exploratory studies suggest correlations between collaboration dimensions, learning gain and haptic mediation; and promisingly, a statistically significant relationship between the number of hCIs and learning gain in two environments. We close with a set of design considerations derived from our thematic analysis.

5.2 Introduction

From early childhood, haptic communication is one of the most intuitive ways that we perceive and interact with our environment, and communicate our experiences and feelings to one another. Haptic information can contribute to rhetorical discourse, in which participants use subjective interpretations of a cue. It can also be used for *dialectical* discourse, by providing partners with factual information from which they can build an objective conversation.

Recently, considerable attention has focused on social haptic touch, both direct and technically mediated. *Socially*, haptic technologies are often used to convey interpersonal information such as emotion or reassurance [95, 216, 232], from a friendly pat on the arm [17, 27] to deeper affective communication [196, 200].

When we consider touch and person-person collaboration more broadly, we

Figure 5.1: The Blind Men and the Elephant. People who have never seen an elephant try to conceptualize it via touch alone, either joining or fractured by their different perspectives [101].

see many cases where *functional* haptic information can be used to support users in jointly carrying out many kinds of tasks. While purposes such as coordination assistance are most common [31], in this work we are interested in how our sense of touch can engage useful collaboration as individuals attempt to perceive characteristics and properties of an object or environment [92, 131].

The theory of *collaborative grounding* offers a base from which to regard a modality's potentiality in the development of joint activity where both content and the process of expressing it matter. Grounding refers to a process by which individuals maintain and/or develop some degree of mutual understanding, effectively finding joint "common ground" from which they can further communicate or ideation [13]. A key marker of collaboration, the process of grounding has been modeled with stages of *monitoring*, *diagnosis* and *repair* [48] (Table 5.3).

Recognition of how touch can figure into collaborative grounding is ancient, traced through Buddhist and Hindu texts to before 1500 BCE [226]. Figure 5.1 illustrates the parable of a group of blind people physically appraising an unfamiliar creature. Each person must express their own perception to help the group conceptualize the elephant [74] (grounding success) – or, as the story is often told, they are unable to resolve their apparently conflicting perceptions (a grounding failure) and confusion ensues. Despite this plus a large body of research on grounding via mediating visual and audio systems, touch is typically overlooked [209].

Our goal in this paper is to study the challenges and the design space of using haptics to attain collaborative grounding through the physical sharing of objective experience. We focus here on functional haptics through force feedback: the simulation of physical attributes of a virtual object. Force feedback devices can convey information such as weight, velocity, collisions, vibrations and attraction/repulsion, through user movements known as *exploratory procedures* [131]. In comparison with direct interactions with physical systems, sensory display of these physical parameters are often absent from graphical and auditory media, and from virtual simulations of phenomena that are distant, large or small, or abstract.

Haptics does bring technical challenges, with accessibility forefront. While visual and auditory media require just a graphical screen or ever lower-cost virtual reality headsets, inexpensive yet high quality force-feedback displays are still rare. Some studies have used high quality haptic displays to assess the potential of haptic feedback on learning [83, 156, 157] with encouraging, if narrowly focused, results; but high cost puts them out of reach for many contexts and for at-scale studies.

'Do-It-Yourself' (DIY) haptic movement [144] suggests that more affordable devices [67, 149] might be feasible for educational purposes, and studies have begun to assess their benefits in different learning activities [44, 155, 201].

As with any collaboration or learning technology, results can be difficult to objectively assess due to the many individual and situational factors in play. Hence, our approach at this stage is to focus on exposing and analyzing collaboration *strategies* in response to different conditions and opportunities, rather than statistical outcomes.

We thereby define two guiding research questions:

1. **RQ1:** How can force feedback affect the process of grounding in collaborative learning environments?

2. **RQ2:** What haptically enabled strategies do learners use to create mutual understanding?

5.2.1 Approach and Contributions

While some studies have examined the use of haptics to assist embodied cognition or as an additional information channel, we are not aware of any that target the use of force cues in a grounding process. We set out to fill this void.

Context: We chose **collaborative learning** as our lens for investigating haptic grounding: specifically, comparing how a range of forms of low-cost haptic force feedback can support grounding as two individuals work jointly to understand a concept. Collaborative grounding is recognized as an insufficiently met need in learning technology [13], and we see its lack in recent experiences with widespread online learning. We posit that appropriate force feedback can positively alter the course of grounding with young learners.

We observe this by analysing how learners use a haptic tool to convince themselves and a partner of an idea, wherein they must provide evidence and consolidate arguments. Rather than statistical evidence of changes in topical understanding, we explore evolution of learners' mutual understanding of the topic and of one another's beliefs.

This framing exposes the strategies and purposes by which learners share information with a partner via functional haptics. The resulting insights into haptic impact on collaborative strategies can also inform single-user learning environments, and help designers recognize opportunities to provide learners with new information and invite them to perform a particular activity or reflect on outcomes.

Platforms: Haptic displays differ in mechanical capabilities and fidelity of haptic rendering [201]; haptic feedback will different roles depending on how the application deploys it [112, 147]. For broad perspective on haptic features key to collaborative learning and to reduce the risk of a single device's limitations masking our findings, we utilized three different devices (Figure 5.2. For each, we designed a unique learning environment which leverages that device's special capabilities.

Two Studies: Our first study used one device, the MagicPen, to confirm that learners were using haptic information during the process of achieving mutual understanding,

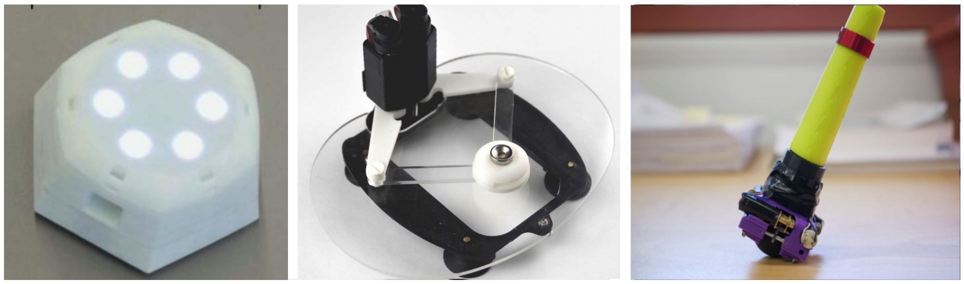

Figure 5.2: Three educational haptic devices used in this work. From left to right: Cellulo [171], Haply [67], and MagicPen[113].

as well as how and why (n=8; 4 dyads). Study 2 utilized three devices/environment combinations (n=24; 12 dyads). We analyzed learners' open dialogues around *haptic critical instances*, based on established dimensions of collaboration [151]. Through pre- and post-tests we looked at early trends in learner's knowledge changes, and linkages to collaboration patterns. Through both studies, we identified conditions that encourage a pivot to haptic use for grounding, and observed strategies by which learners used haptics in the different environments. We organized findings into four thematic categories.

We contribute, encapsulated as reflections and design implications in our Discussion:

- Insights into the utility potential of force feedback in collaborative grounding, as enabled by **device capability**;
- Identification of **strategies** whereby learners were effective in using force feedback for grounding as well as knowledge confirmation and learning;
- Articulation of the **layered roles** that users tend to assign to the haptic modality as they collaborate, from communication mediator to exploration, design and learning tool.

5.3 Related Work

5.3.1 Educational Haptics

Educational haptic platforms exploit both visual and haptic cues to create interactive environments that teach abstract topics in physics, math, and other fields of science. They provide the tangibility of digital manipulatives, and can combine force feedback with visual cues to create a more compelling experience in STEM (Science, Technology, Engineering, and Mathematics) learning scenarios. Previous work has demonstrated the benefits of physical manipulatives in education, particularly science education. While some studies suggest physical manipulatives directly impact learning, resulting in higher test scores when students use these manipulatives [199], other studies highlight indirect factors that can lead to learning gains such as increased collaboration [234] or new possibilities for classroom orchestration offering new possibilities to teachers to better orchestrate their classroom, *e.g.*, remote teaching and learning [51].

Haptic devices, like other physical manipulatives, add physicality to abstract concepts [44]. As a result, in the learning environments, learners can reason about the movements and behaviour of different objects, which may be otherwise inaccessible. Moreover, haptic devices exploit the sense of touch to provide an additional channel of information for students. Through haptics, learners can feel different mechanical properties as well as perceive the shape, volume and weight of objects. This non-verbal channel may help build intuitive understanding of both symbolic and iconic concepts in the learning scenarios [44, 112].

5.3.2 Haptic Communication

Haptic communication can fall into dialectical or rhetorical discourses, depending on the type of information provided (objective/factual vs. subjective/emotional). The application of educational haptics needs to be objective in order to meet intended needs, *i.e.*, conveying specific information or creating a specific experience. As a result we expect that teachers, students and hapticians who are designing these interactions have similar interpretations of their haptic perception during the learning activities.

A large body of work explores the role of sensory feedback in science learning. The importance of the sense of touch while interacting with learning environments invites two lines of research. The first investigates the significance of physicality in physical manipulatives in comparison with learners' interactions with virtual manipulatives [45, 123, 236]. The second, and more aligned with the focus of this research, studies the educational benefits of embedding haptics in virtual labs (simulators). The majority of these empirical examinations test the added value of force feedback to the virtual environment through a quasi-experiment with at least two conditions (with and without presence of haptics) [85, 225]. A majority of the activities, however, find a learning gain for both conditions with no significant differences between them [112, 234]. Educational haptic literature tends to focus on quantitative score changes to measure learning, but fails to establish a framework for *What*, *How* and *Why* force feedback should be employed in learning environments [234].

Here, we study haptics in a collaborative context. We will investigate how haptic-enabled devices impact collaborative discourse in different learning scenarios. This research attempts to find a collaborative benefit to educational haptics rather than a causal relationship between haptic use and test scores.

5.3.3 Collaborative Learning

The demonstrated benefits of collaborative learning (where individuals jointly try to construct, negotiate and share understanding) [28, 52]) include learners' planning more accurately and generating more solutions, as well as boosting individual performance in near-transfer problems [14]. Collaborative test-taking demonstrates a significant increase in test score and students' retention of learning materials [23]. Supportive tools can create more opportunities for collaboration and significantly impacts collective knowledge building dynamics and consequent individual learning [238].

While many collaborative learning technologies exploit the auditory and visual channels to mediate peer interactions between peers, the importance of the sense of touch as a medium for grounding is largely unexplored. In focusing on the potential of deploying touch to support grounding, we first require a haptic platform

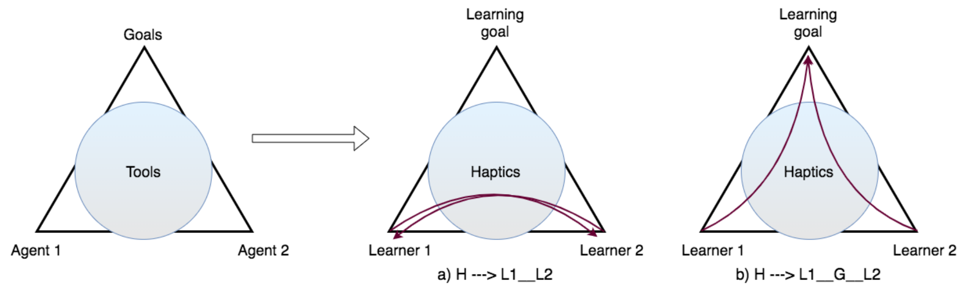

Figure 5.3: A modified framework based on Baker *et al* [13]. (Left) A simplification of Baker's original model for analysing the effect of a tool on grounding. (Right) In our work, the tool is Haptics, Agents are Learners, and we focus on two impacts (red arrows): a) Learners' haptics-mediated mutual understanding of each other; b) Learners' haptics-mediated understanding of the goal.

and accompanying learning environment that intrinsically encourages students to scaffold one other's learning. However, in a chicken-and-egg quandary, we know little of *what* tool and learning environment properties will best accomplish this.

Baker *et al* present a simple framework (Figure 5.3, Left) for understanding and analyzing the role of tool-mediated grounding in mutual understanding in a collaborative learning task [13]. We build on this framework to conceptualize our work.

5.4 Framework and Tools – Haptic Devices and Learning Environments

5.4.1 Framework for Exploration of Haptic Grounding

We build on Baker *et al*.'s framework [13] to explore how introducing a haptic modality (force feedback) can influence different aspects of grounding. Our primary focus here is on how haptics might help peers to achieve grounding and to share their experiences and thoughts across different learning environments, rather than whether it can improve learning by an individual (Figure 5.3, Right; and detailed below).

Agents: Our agents are two learners expected to have a meaningful mental state –

able to act and interact with other learners and the haptic tool – and the ability to infer the state model held by their partner.

Goals: We are studying three learning scenarios, each of which has several defined learning goals. Learners must collaborate to perform each assigned task.

Tools: We use a haptic device paired with a custom virtual-environment learning environment (described next), which together mediate learners' collaboration as they jointly perform certain tasks

We expect to see: Among *entities*, we assume that goals and tools are static, and focus on cognitive changes within our agents (learners). In terms of *relations*, we focus on the changes caused by haptics in the agents' mutual understanding of the goals and the mutual understanding that develops between them. Changes in each entity or relation can impact others and could be bidirectional but are beyond our scope.

5.4.2 Haptic Devices

We chose three low-cost haptic devices that met requirements of being available (both as physical devices and in programmable access), sufficiently robust for this study, and as a group representing an important and diverse range of forms among the broad category of manual, 2-D force-feedback devices (Figures 5.2 for device closeups, and 5.4 in situ).

All three devices require no calibration (*i.e.*, in an experimental context are 'walk up and use') and have been used in published studies.

We paired each device with its own custom collaborative learning environment and activity, to best exploit its capabilities (Table 5.2). While theoretically possible to implement a full factorial comparison of all environment/activities with all devices, we were dissuaded by concerns of study size, participant fatigue, poor outcome targeting, and inconclusive results due to poor haptic quality from inherently unsuitable pairings. Conversely, devising a single environment that was "fair" on all three devices would mean sinking to a lowest common denominator of engagement and conceptual breadth.

The device-specific attributes we sought to exploit are as follows.

Cellulo [171] is an untethered puck that moves relatively slowly but autonomously, localizing on watermarked paper, and pushes gently against the user's hand. Its accurate localization and independent robot movement enable collocated interactions with visual cues on a printed poster paper in the **Pressure Lab**.

Haply [67] is a DIY version of a pantograph, one of the most common planar mechanisms, with a compact constrained workspace and reasonable responsiveness. It has probably the least familiar form for students. Its fast rendering and high quality/strong motors gave the best fidelity of the 3 systems, and were crucial for rendering object impacts in the **Collision Lab**.

MagicPen [113] is an untethered stylus which provides forces generated through a frictional rolling contact on an arbitrary surface; it needs assistant to stand up, but can actively drive a user's hand around and display a variety of force environments [116] in a theoretically unlimited workspace. MagicPen has the lowest expected learning effort, since it mimics pens and styluses, which most students in North American schools are familiar with. Its fast rendering and response speed help learners to feel the immediate changes in the force fields needed for the **Electrostatic Lab**.

Haply and MagicPen's environments are developed in Processing (Java) and participants interact with them with a laptop and a graphical screen. The Cellulo environment is implemented in QT (C++/JavaScript), a cross-platform toolkit that connects the tangible robots and printed environment wirelessly to a tablet.

Participant dyads each had their own Cellulo and MagicPen, but shared a Haply: in the Collision & Momentum Lab, "throwing" objects (so a partner could feel the impact) worked better on a trackpad than with Haply's constrained workspace.

One might question the reason behind using these three haptic displays instead of high-quality haptic devices such as Phantom Omni. To answer this question we should highlight the special needs and unique features of each haptic device besides the affordability constraints.

We needed an overllapable workspace for the Electrostatic Lab with MagicPen study (untethered). Similarly, Collision Momentum Lab demands a 2D haptic device since forces in 3D could make the experiment very confusing. We also needed a large workspace for the Electrostatic Lab which was beyond the workspace of the Phantom Omni haptic display.

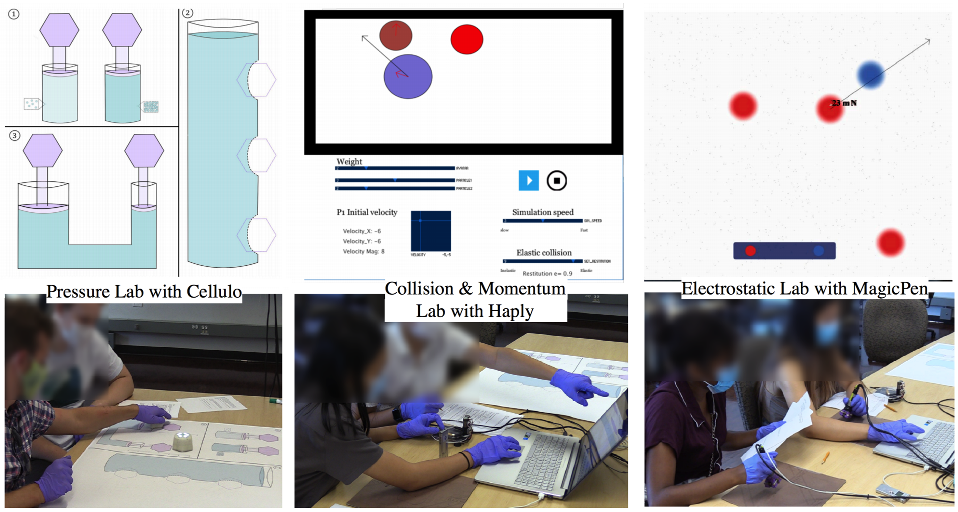

Figure 5.4: Three virtual environments paired with haptic devices. From left to right: Pressure Lab with two Cellulos, Collision and Momentum Lab with One Haply, and Electrostatic Lab with two . Gloves and masks were used as part of our institutionally approved COVID-19 safety protocol.

5.4.3 Haptic Environments and Associated Learning Tasks

Our haptic environments are shown in Figure 5.4 (screen or poster shots and in-situ views). We developed learning goals in reference to the local region's high school science curriculum. for relevance to high school students (year depending on curriculum component). We designed each learning activity with attention to minimizing the interface's cognitive load during learning [13].

Pressure (Fluids and Hydraulics) – Paired with Cellulo

Learning Goals (Grade 11:) We address natural and constructed fluid systems, specifically how compression of fluids can be used to power these systems. Upon completion, students should be able to observe evidence of pressure, and measure how fluids react to pressure and model fluid systems [168].

Haptics Role: We focus on conceptual understanding and applications of fluid properties of compressibility, flow and hydraulic systems. These can be difficult to illustrate in the classroom, but Cellulo can bring these concepts to life, and allow the learner to interactively explore and feel the fluid's behaviour. Through its movement Cellulo demonstrates the behaviour of the liquid coming out of a bucket of water or the pistons on a hydraulic jack so that students be able to feel the pressure on each scenario.

Learning Task: Learners first perceive the compressibility of a gas (*Familiarization*); then in pairs, the effect of pressure on flow and of hydraulic multiplication (*Accumulation*). Finally, they must parameterize a hydraulic jack that can lift up a car (*Design*), deciding whether to put the car on the larger or smaller surface area of the jack and to fill the jack with air or liquid.

Collision and Momentum – Paired with Haply

Learning (Grade 8): In Newtonian mechanics, linear momentum, translational momentum, or simply momentum is the product of the mass and velocity of an object. It is a vector quantity, possessing a magnitude and a direction in three-dimensional space. Upon completion, students should be able to analyze, evaluate and apply information about the conservation of momentum in closed systems and collisions.

Haptics Role: In an isolated environment, the momentum of the objects before and after the collision are equal. The Haply renders the law of conservation of momentum in three different linear interactions: collision where objects do not stick together; collision where objects do not stick together; and explosion. The Haply renders an impulse force on the avatar object "held" by the user when it is "thrown" by the partner using the trackpad.

Learning Task: If m is an object's mass and v is the velocity (also a vector), then the momentum is $P = mv$. We ask learners to change mass, velocity and coefficient of restitution and feel the impact on the momentum (Familiarization and Accumulation). They also have to set the parameters in an environment where two particles collide and stick together (Design).

Electrostatic Lab – Paired with MagicPen

Learning Goals (Grade 8): We address Coulomb's law, which states that the force between two charged objects is proportional to the magnitude of the charges and inversely proportional to distance between them. Upon completion, students should be able to predict electrostatic dynamics between charged objects in one dimension. They should also be able to explain the relationship between electrostatic force and distance between charges.

Haptics Role: MagicPen renders electrostatic forces as linearly proportional to the amount of charge, and follows the inverse square relation to the distance [114]. Generated forces from point charges are small in real life, and need to be scaled up to be perceptible [148]. MagicPen demonstrates the difference between attractive and repulsive electrostatic forces by drawing the two partners' devices together in the former case and pushing them apart in the latter.

Learning Task: Learners are asked to construct a physics example, in order to experience Coulombs law and governing attractions between point charges (Familiarization). They can construct the model and move the placed charges using the graphical interface (Design), and can control the main charge with the visual information about the force and test their hypothesis (Accumulation).

We describe data collection processes associated with both of the studies outlined in Section 5.2.1 and further detailed in Sections 5.5-5.6 respectively.

5.4.4 Dyad Participants

Because collaborative learning requires communication with a peer, both of our studies were built around dyad participation, as have others [44]. In contrast to a solo think-aloud approach [112, 148, 184], pair work is intrinsic to collaborative learning, and induces pairs to speak to one another, providing a window into their thinking as well as verbal and nonverbal cues regarding their orientation and emotions towards their partner.

Demographic: To balance representation of high-school-age learners with ethics and safety constraints during the pandemic, we designed both studies for early-stage university undergraduate students who had never taken a university-level Physics classes, chosen to simulate target levels of Physics knowledge.

Recruitment: Due to the COVID-19 conditions in place for both studies, we asked each recruited participant to bring a partner so that each dyad was composed of individuals who were within the same social "bubble" – partners could then sit closely as they worked. This type of pair recruitment simulates secondary and

Table 5.1: Summary of four-part experiment procedure for a single device/environment combinations, for Study 1 (MagicPen) and 2 (all three combinations).

Stage	Description	Example
1. Pre-Test	Brief assessment of existing conceptual understanding	What is an electrostatic force? What interactions exist between positive, negative, and neutrally charged objects?
2. Activity (a): Familiarization Exercises	Dyads answer questions on paper relating to physics concept	Will these two same charge balls swing apart or together?
2. Activity (b): Accumulation Exercises	Predicting and reasoning about the answers. Using the haptic device to do the experiment and compare it to their prediction.	Consider the following charge formation. If the negative charges are held fixed, explain the behaviour of the positive charge in terms of the forces between the charges.
2. Activity (c): Design Task	Individual, then as dyad, design a scenario that could be tested using the learning environment	Arrange charged particles so that net force on a target point is neutral "How did you convince your partner that your design worked?"
3. Post-Test	Re-assessment of conceptual understanding	Similar to pre-test
4. Usability & Perception Survey	Determine enjoyment of activity and perceived usability of device	Likert scale for enjoyment, usability, and usefulness

primary school experiences in which students typically know each other prior to collaboration, and may be important in terms of comfort level in having frank conversational exchanges.

5.4.5 Core Experiment Procedure: One Device/Environment Block

For each device/learning activity block (executed once in Study 1 and three times in Study 2) we conducted a four-stage data collection procedure, following [148]: (1) pre-test to measure prior learner understanding and knowledge; (2) primary collaborative activities to observe the interaction between learners and possible effect of force feedback on the quality of grounding; (3) post-test activity to measure

Table 5.2: Environments paired with haptic devices

Device	Environment	Mechanism	Key Device Features	Software
Cellulo [171]	Pressure Lab	A pocket-sized handheld mobile robot with the ability to generate force feedback (2D+roll). It uses and embedded camera for absolute global localization	Accurate localization, large workspace, holonomic motion, Force Amp: 1.75N	QT (C++)
Haply [67]	Collision & Momentum Lab	Uses a 2-degree-of-freedom pantograph mechanism with parallel joints and optimized dynamic performance	Fast rendering >1kHz, High quality motors, Force Amp: 4N	Processing (Java)
MagicPen [113]	Electrostatic Lab	Friction of a rolling ball over a surface creates 2D force feedback to the user's hand	Fast rendering >1kHz, fast response speed, large workspace, Force Amp: 1N	Processing (Java)

learning gain; and (4) usability and perception survey to reflect on the learners' interaction experience with each device. Table 5.1 summarizes the experiment procedure.

Stages 1 and 3: Pre- and Post-Tests Before and After Learning Activity

Participants completed pre- and post-tests individually and without access to the haptic devices, to detangle individual understanding of the physics concepts, and determine whether they were able to apply mutual concept understanding to individual work. Pre- and post-tests consisted of similar 5 questions, designed to capture intuitive conceptual understanding (3 questions), and knowledge applications (no calculations) (2 questions). Questions were multiple choice and short answers and administered on a laptop.

Stage 2: Three-Phase Learning Activity

Learner common ground and knowledge accumulation are both important for collaborative discourse [35]. We designed our learning activity to ensure that participants could achieve a common ground (Familiarize), and then add to it (Accumulate), both steps essential for our analysis. We added a third phase (Design) so we could assess quality of collaboration along the scale of minimal to optimal effort [48].

(a) Familiarization: Researchers familiarize learners with the haptic device and

environment, then assess learners' mutual understanding of the basic concepts covered in the present lesson, as a pair.

(b) Accumulation: Learners add to their common ground by working together to predict answers and solve problems through experimentation within the environment. Our study goal here is to understand how learners use haptic information to fill in gaps and add to mutual understanding; to do so, we must motivate them to construct examples that their partner can experience.

(c) Design: Working on paper, learners design a scenario that could be tested using the learning environment. Then, they collaboratively discuss and test each of their designs, and decide which meets the requested criteria best (and why). The design phase offers learners creative control. Differences in design often lead to constructive debate between partners which furthers their grounding. Ideally, learners then *repair* those differences [48].

Stage 4: Usability and Perception Survey

We conducted a survey after each activity asking learners about the quality of haptic perception and ease of interaction with environment and device [202], consisting of four Likert questions (1-strongly disagree to 7-strongly agree) on (1) enjoyment of learning physics, (2) ease of interaction with the device, (3) ease of interpreting the force feedback, and (4) usefulness of device in learning physics.

5.4.6 Full Procedure

The experimental setup consisted of a laptop and haptic devices (Figure 5.4. Environments – one (Study 1) or three (Study 2) – were arranged in nearby stations. Following a welcome and consent finalization process, learners were seated together and shared a monitor at the current station, then commenced the environment/device blocks.

We required learners to wear masks and gloves for COVID-19 safety, and took other sanitizing precautions before and after experiments. Approval for both studies was obtained from the University Behavioural Research Ethics Board.

5.4.7 Collected Data

The above procedure resulted in the following data:

- Individual demographics, including physics expertise and intra-dyad relationships
- Pre-test and post-test scores
- Learning activity written responses
- Voice and screen recordings of dyads completing the learning activity
- Videos of the participants hands operating the haptic device during the learning activity
- Log of force, velocity, and displacement of the haptic device (Only study 2)
- Usability questionnaire

5.5 Study 1: Confirming the use of haptics in grounding

In our first study (based on Electrostatic Lab paired with MagicPen), we looked for instances of and opportunities for using haptics in collaboration. How and the degree to which learners incorporated haptics into explanations and grounding strategies would direct our further inquiries. Specifically, we considered:

1. At what *stage* of grounding did learners tend to incorporate haptics?
2. What haptic *gestures* were made with the device?
3. What were learners' *intentions* with haptic information (*e.g.*, introduce or revisit content, evaluation,)?

5.5.1 Method

Procedure:

For data collection, we followed the general procedure of Table 5.1.

Participants:

We recruited four participant dyads (n=8) through word-of-mouth[2]. Two dyads described their relationship as "close friends", one as "family members" and one as "casual friend". Seven participants were female; all were within 16-24 years of

[2] Study 1 occurred days before the first university closure due to COVID-19 in March 2020.

age with the majority between 17-19. Each participant was compensated with $15 (actual session duration 45-60 minutes, M=52).

5.5.2 Analytical Approach

We developed a three-layer qualitative qualitative coding approach to find collaborative actions based on screen capture videos and audio recordings, then successively unpacked them for haptic import.

First, we identified *what* (collaborative grounding acts were executed through any modality, and as in the conversational analysis method of Porcheron *et al* [182], we used these boundaries to fragment the data into collaborative learning episodes, with timestamp-maintained interaction structure. This fragmentation was necessary to discover and report co-occurrence of actions, intentions and strategies for use of the MagicPen.

We then looked for *how* (gestures made with the haptic devices and associated with those grounding acts, hereafter referred to as *haptic gestures*); and finally *why* (the intention behind the haptic gestures). For these, we developed a haptic gesture coding system inspired by Yohanan *et al* [232]: we used grounding acts identified in the data logs to segment our qualitative data then categorized haptic gestures used in each grounding act, and the intentions behind them.

Identifying Grounding Acts (What):

We first had to identify all grounding acts and stages (monitoring, diagnosis, and repair) (Table 5.3 [48]) as they occurred.

Collaborators *monitor* each other when they determine the information and beliefs that their partner has. This can be done actively by expressing one's own opinion, for example: *[P8] "So these arrows are going like this."*. *Diagnosis* is explicitly acknowledging a difference in belief or information access between collaborators or discrepancy between their beliefs and new information: *[P7] "No, no, no, no. I think we're wrong here."*. A grounding act is complete once collaborators have gone through the stage of *repair*, or adjusting their beliefs to match so that they become part of the common ground. *[P8] "It will be pulled down."— P7: "Ok."*.

Table 5.3: Grounding Acts, repeated from "Grounding in Multi-Modal Task-Oriented Collaboration" [48]. 151 grounding acts were identified in Study 1.

Acts	Examples
Monitoring (53%)	A infers that B accesses X A infers that B notices X A infers that B understands X A infers that B (dis)agrees with X
Diagnosing (25%)	A joins B to initiate co-presence A asks B to acknowledge X A asks B a question about X A asks B to agree about X
Repairing (22%)	A makes X accessible to B B communicates X to A A repeats-rephrases-explains X A argues about X

Table 5.4: Haptic Gestures made with the MagicPen. 130 gestures were identified in Study 1.

Gestures	Haptic interaction examples
Construction (23%)	Adding, deleting, moving points of interest (POIs)
Physical Collaboration (23%)	Switching pens, switching positions, moving location requested by partners, asking partner to move their avatar
Exploration (42%)	Moving towards/away from POI, varying speed, exploring environment edge, scanning environment, detailed inspection, bumping into POI
Play (12%)	Chasing partner's avatar, bumping into POI repeatedly, circle a POI

For each dyad, we identified grounding acts based on verbal statements and intonations from voice recordings. The above quotes, all drawn from the same 40-second period, illustrate a common sequence of monitor, diagnosis, repair.

Identifying Haptic Gestures (How):

To determine whether the MagicPen was being used for collaborative acts, we looked at **how** learners were using the device. We first identified all gestures and action occurrences that participants made with the MagicPen (a total of 130 events). Then, we classified these occurrences into four categories: *physical collaboration*, *construction*, *exploration* and *play* (Table 5.4). This categorization was based solely on hand movements and the resulting on-screen avatar movement, not on other clues such as utterances or the task at hand.

Identifying Intention (Why):

Finally, we surmised intent behind each haptic gesture using recordings of participants' onscreen avatar movements, voice and hands. Unlike the gesture coding

Table 5.5: Presumed intention behind each haptic gesture (adapted from Cesareni's "Global Conversational Functions" [30]). 174 presumed intentions were identified in Study 1. Asterisks denote functions or gestures we have added to Cesareni's table to cover all actions demonstrated by our participants.

Intention	Example
Introduce new information (32%)	Introduce personal ideas, Introduce information from a reliable source, Introduce examples drawn from experience, Pose research question or problem
Revise previous information (19%)	Elaborate own ideas, Elaborate other's ideas, Synthesize, Repeat other's ideas, Repeat own contribution, Respond to other's ideas*, Resolve an ambiguity*, Understand partner's perspective*
Evaluate or reflect (33%)	Express meta-cognitive reflection, Comment, Evaluate (Reason*), Test their hypothesis about the environment*, Determine magnitude/limitation of haptic perception*, Share experience with partner*, Take turns with roles/positions*
Maintain relationships (3%)	Express agreement/disagreement, Maintain social relations, Convince partner of a hypothesis
Fun* (8%)	For fun, To enjoy the session
Logistics* (5%)	Request an action from a partner, Confirm completion of a task, Determine the next task

which focused on how users relied on the device, here we unpacked the goal that participants were working towards with each haptic gesture, which often depended on the activity task or their partner interactions. Although this process is imprecise, Cesareni [30] identified four broad types of utterances that support knowledge building: *introducing new content*, *revisiting content*, *evaluation*, and *maintaining interpersonal relationships*. Some haptic gestures did not seem to build knowledge, so we added two more categories: *play* and *logistical coordination*. These six categories covered the likely intent of all the MagicPen gestures we observed (Table 5.5).

5.5.3 Results

Two coders collaboratively discussed and coded the script, using MAXQDA software [126]. Grounding acts and haptic gestures were relatively straightforward to observe; intentions coding benefited from inter-coder discussion and calibration.

We found a total of 455 codes (M=114 per group, and SD=47) including: 151

grounding acts, 130 haptic gestures, and 174 presumed intentions.

In studying co-occurrence, we first compared each code with others outside its coding family. For example, for each intention behind haptic use (*why*), we looked for frequent co-occurrence with all haptic gestures (*how*) and collaborative act (*what*) codes. These co-occurrence counts were based on frequency within collaborative learning episodes. As a result, co-occurring codes sometimes came from both participants, which we deemed admissible since both people were contributing to the problem solution, the grounding process and the haptic experience (by interacting with each other's avatars).

The result (Figure 5.5) suggests that learners tended towards recurring *x-y-z* patterns – accomplishing a specific grounding act *x* by using the device in manner *y* with the intention of expressing *z*. We elaborate on patterns with highest co-occurrence (large circles).

Grounding act (Monitoring) + Haptic gesture (Exploration of environment)

(31 co-occurrences). Learners often used the MagicPen to determine something about a layout of point charges that they had previously created. During the subsequent exchange, they used force feedback from a fixed point charge to form a belief about the strength of a force they experienced – for example:

> *[P7] Oh, he doesn't want to stay close to this blue buddy.*
> *[p8] Oh but maybe my guy needs to stay close. You just stay-*
> *[p7] No, see he's trying to move away.* [P7 moving the pen towards the blue]
> *[p8] Other repulsive charges are stronger than your attraction to me.*

Haptic gesture (Exploration) + Intention (Evaluate/Reflect)

(32 co-occurrences). Perception of changes in forces' amplitudes and direction enabled participants to evaluate their hypothesis and reflect on their conclusions. They searched device behaviour for evidence that confirmed or denied their hypotheses about point charges. This pattern, which appeared across all the dyads, suggests that they were employing haptic information in their learning process, to test their intuitive understanding of point-charges.

Grounding act (Monitoring) + Intention (Introducing new information) + Haptic gesture (Construction)

(26 co-occurrences). Participants often shared new information during the Monitoring stage. In most cases this included using the MagicPen to create new point charge layouts while verbally sharing a hypothesis, sometimes then relying on the force feedback to confirm or reject their hypothesis (16 of the co-occurrences).

> *[P5] Okay. First let's line them up like this one. Like this one and now, now what we do... We just let it go.*
> *[P6] Yeah.*
> *[P5] Seems like they... we'll swing away from each other due to the repulsive electrostatic force between them.*

Taken together, these two processes of introducing information and exploration allowed learners to first create a point-charge system and then use it to test their hypothesis.

Sequence of coded events

We used a code map to plot code overlaps by including a transcript paragraph before and after each episode.

In Figure 5.5b, each circle denotes a coding subcategory, which cluster into three groups. We used the MAXQDA software's clustering tool and increased the number of clusters (K) until we arrived at a plateau. From K=4 onward, we observed appearance of clusters with a single member.

Cluster 1: Yellow: We see strong overlap between *Construction*, *Physical Collaboration* and *Introduction* codes, which manifested as participants using their devices to jointly build an environment. This was often to rapidly test an environment they were asked to explore or a hypothetical environment. For example, learners might ask their partner to move their avatar to experience the effect of this movement on their own avatar. Introducing new information was often followed by monitoring acts where participants tried to ensure that their peer understood the new topic.

Cluster 2: Pink: Learners usually employed their haptic devices to explore the environment and then evaluate and reflect on their hypothesis. While we expected

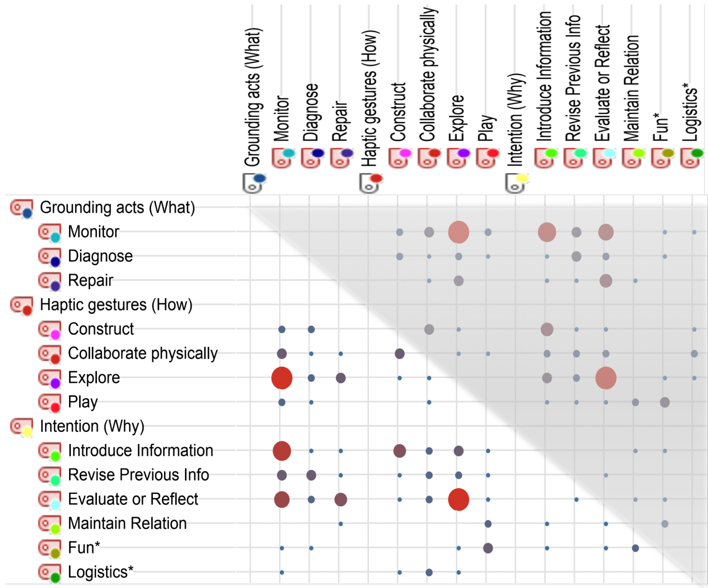

(a) Co-Occurrence of Codes. Circle size denotes frequency of co-occurrence.

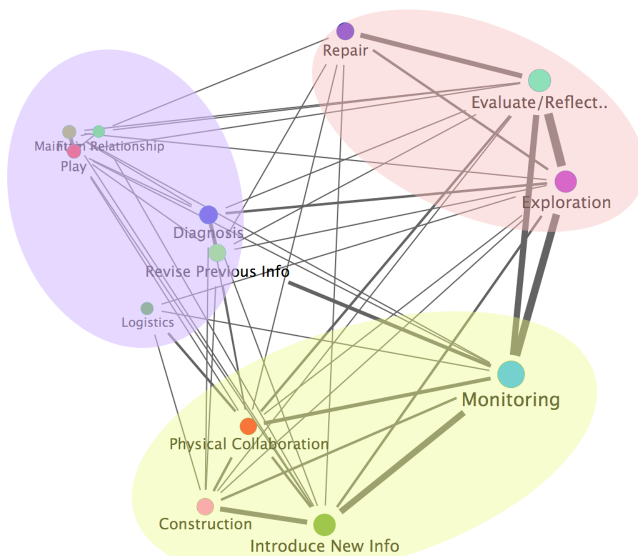

(b) A code map plots codes according to similarity. Circle size denotes frequency of code assignments; intercode distance represents overlap on the same script. Connecting lines indicates which codes overlap or co-occur (within a paragraph before and after). Line thickness denotes code co-occurrence.

Figure 5.5: Coding relationship analysis (Study 1).

that users would talk about their haptic experience to ensure comparable experiences, we saw that some preferred to first try and reflect on the environment by themselves. Hypothesis evaluation often lead to a repair (consensus) act, in which they might rely on the haptic experience to convince their partner. We saw that haptic perception could be personal, with individuals preferring to discuss higher level information; however, when this higher level failed, they would restart by agreeing that they were feeling the same behaviour. This similar haptic experience would eventually create a base from which they could build to reach to a common conclusion.

This can be seen in Figure 5.5, where code closeness implies overlap in use, but also sequence. *Monitoring* is closer to *Exploration* than to *Revising*, which suggests that monitoring is often followed by or precedes exploration. The thick connection indicates a high frequency of this co-occurrence.

Cluster 3: Purple: In *Monitoring*, learners tended to discuss their ideas and revise or repair them. This often occurred through conversation rather than using the haptic device, potentially because the information required to repair a rift between two beliefs was more conceptual than experimental. As such, learners may have found it easier to express this information verbally than haptically.

Learning Gain and Usability:

We measured learning gain by subtracting the individual's post-test score from their pre-test score (the assessment questions are presented in Appendix C). We observed a positive learning gain (M=1.125 SD=1.36 out of 5 points). The results of the user experience survey showed reasonably good scores (at least 5 out of 7 Likert scale points in all dimensions). These were: Enjoyed learning (M=5.62, SD=1.30), Easy interaction with device (M=5.37, SD=0.74), Understood the applied force (M=5.38, SD=1.06), and Useful for learning physics (M=6.25, SD=0.89).

5.6 Study 2: Haptic Critical Instance Analysis

Study 1 suggests the utility of haptic interaction during collaboration (learning gain), revealed examples of how participants naturally use haptic gestures in their conversation (RQ1), and provided insights on the strategies of (when, what and how) to use haptics for grounding (RQ2). However, we could not tell when participants were relying on haptic or visual cues, and could not search for correlations between learning gain and haptic use.

In Study 2, we utilize *haptic critical incidents* to understand how different forms of force feedback influence collaborative learning, based on three environment-device pairs.

5.6.1 Full Study Design and Protocol Modifications from Study 1

Design and Sample Size:

We ran a within-subject dyad study with 2-hour sessions that produced comparable transcript and videos across three haptic educational devices. We chose the single-session format to minimize participant transitions under COVID-19 safety protocols.

In choosing sample size, we considered our chief goal of digging deeper into moments where haptics plays a key role in the conversation or collaborative actions, and sought sufficient number and diversity of these moments to study rather than statistical significance [106]. Each dyad-session produces many. Thus, we sized the study based on two counterbalanced repeats of the three device-environment pairs: 2 repeats x 6 orderings, leading to n=24 (12 dyad-sessions).

Procedure:

The data collection procedure was similar to that of Study 1, except that the basic 4-step elements (Table 5.1) were repeated three times per session, once for each device/environment combination, order of combinations counterbalanced across dyad-sessions. After finishing the activity in each environments, learners took a 5 minutes break and then we invited them to start the next activity.

Participants:

We recruited 24 participants through online advertising among first or second year non-STEM university students. Recruiting was done by pairs, as for Study 1. We compensated each participant $40 for their participation in a 2-hour session (actual duration 95-120 minutes, mean 102).

5.6.2 Additional Data Collected

To Study 1 data (5.4.7) we added a debrief interview after each activity. We asked students to reflect on advances made by the dyad and any disagreement. This interview was similar to the debrief interview after Study 1, with the addition of asking participants to identify moments when their understanding of the concept changed. Although participants were not typically able to pinpoint their shift in understanding, this question invited discussion about changes in understanding.

We also created a log of force, velocity, and displacement for our three learning environments. We expected learners to experience different behaviour of the force feedback and perceived changes in magnitude and direction of the force feedback.

5.6.3 Analysing Haptic Critical Instances (hCIs)

Study 2 focus was on places in the activities where haptic use was taking place. As in Flanagan *et al* [60], we used the force feedback data stream to identify moments when students are actively using the haptic device to communicate, play or explore: when device actuators were active, we inferred that the holder was using it in some way. At these points, which we term *haptic critical instances*, we examined all of our synchronized data records, particularly voice recordings and video of participants' hands.

Definition of an hCI: We defined an hCI as any noticeable change in magnitude and direction of force feedback which could be the result of, co-occur with or lead to an utterance, action or movement, so that all relate closely to the aim of the learning activity. We looked at a window of 1 minute before and after the triggering force change. Auto-detected hCIs which did not meet this criteria of connection to other behavior were rejected.

Computational process for finding hCIs: We needed to identify hCIs based on an estimate of *force variability (FV)*. Previous works used *cosine similarity* to capture variability of the force vector [170, 205]; however, this does not offer the temporal localization that we required.

Instead, we computed FV as the interquartile range of the measured force amplitude and direction, within a rolling window on a force data stream. This statistic captures dispersion of force during a haptic interaction gesture.

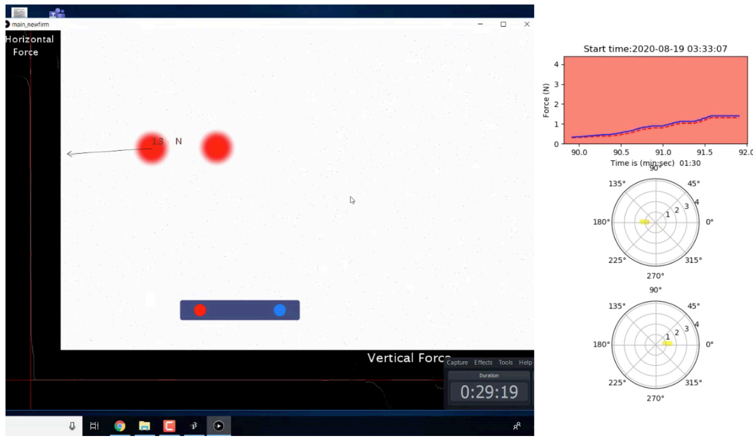

Figure 5.6: Coding haptic critical incidents: **Left:** Coder's view of screen recordings of the haptic environments when the learners are performing the experiments **Right:** Real-time analytical view of force behaviour during any given time in the screen recording. The background of the force amplitude plot (upper right) turns pink to indicate a critical instance, identified from the force logs during playback of the screen recording of a session. This notifies the coders to make and code the critical moment and analyse learner behaviour around this timestamp.

However, we anticipated that the appropriate window size for FV computation would vary by environment. For example, for Collision Lab, relevant changes occur in milliseconds, but in Pressure Lab we would need a window of at least a full second. To identify optimal window duration so as to segment useful chunks of hCI type activity, we referenced frequency domain indicators of typical activity. Specifically, we ran a fast Fourier transform on samples of force logs on each haptic interaction to determine the key haptic frequencies (Figure 5.7c). It should be noted that these frequencies indicate the changes of forces due to the user's behaviour, and the force rendering frequency is at least 10 times higher. Frequencies thus found for each device informed FV window duration, by device/environment.

Automation and data stream fusion: We wrote a script to automatically detect hCIs as an objective indicator of using haptics by each dyad. The algorithm used a 1-D peak finder filter to search for local maximums and minimums, with FV window length set as described above, and was applied to the study force logs to generate hCI time stamps.

From this point, we combined the audio, video and force log data to create a single video-format resource for the analyst. Text transcripts were auto-generated, manually verified, and later added to MAXQDA to facilitate coding.

Manual data annotation: A human analyst team reviewed the compiled resource at auto-marked hCIs locations for all dyad records, to understand the context of each hCI (Figure 5.6), look for events of interest and apply annotations accordingly. These annotations were independent of the 9-dimension collaboration rating described below.

The analysis team consisted of three coders. All reviewed and annotated all 12 dyad records. Annotations were generated across the first 4 records, then subsequently calibrated and refined through discussion during check-ins after each 4 records.

Thematic analysis: We performed thematic analyses for the full set of videos in regions around auto-identified hCIs, using the process detailed by Braun *et al* [26]. We used the 9 collaboration dimensions (column 1 of Table 5.6) as codes [151] by rating the full activity sequence (on a scale of -3 to 3) by collectively considering all the hCIs for that device/environment block. Raters also considered the annotations made in the previous step by themselves.

Specifically, each rater (of three) produced one rating/dyad for each dimension for each of the three study blocks, for a total of 3x12=36 ratings per dimension for each device/environment combination. Throughout, we tried to establish the following information around the hCIs:

- Describe the situation around the haptic usage
- What did the participant/s do with the haptic device or say?
- Why this usage was particularly effective (or ineffective)?

The coders continued this line of questioning as they perused the data, documented supporting details. We continued collecting data until new incidents being analyzed provide few or no additional critical grounding acts. For more complicated activities we checked to examine a greater number of incidents.

We obtained inter-coder reliability of Cronbach's Alpha 0.802, suggesting 'good' agreement.

5.6.4 Quantitative Results

User Experience:

We report general observations on particular affordances that the device/environment pairings seemed to reveal, with no intention of relative ranking; the systems are too different to thus compare.

The results of user experience for the devices/environment combinations are shown in Figure 5.7a. While all devices are rated relatively highly on four dimensions, some outstanding features of each device lead to a slight differences in the ratings. Participant enjoyed learning with Cellulo; as a sometimes-autonomous robot, it was especially engaging and fun to play with. They found Haply easy to interact with, possibly due to the passive nature of the haptic experience in the Collision environment. They found the haptic information from the MagicPen easiest to understand and useful for learning, as they could actively manipulate and explore the Electrostatic environment.

Learning Gains:

We measured learning through pre-tests and post-tests. Electrostatic Lab (MagicPen) has the highest learning gain (M=1.13, SD=1.25 out of 5 questions) followed by Pressure Lab (Cellulo) (M=0.30, SD=1.40) and then Collision/Momentum (Haply) (M=0.02, SD=1.05). We categorized dyads based on their learning gain into three baskets as follow (see Figure 5.7b: High [2-5], Low [0-2] and No learning gain [=<0].

Dimensional Ratings and Correlations:

In total, our automated process identified 2361 hCIs over the 12 dyads' records, of which our team analyzed 1246. Collision & Momentum Lab has the highest average hCI per dyads (M=56,SD=17) then comes Electrostatic Lab with (M=31, SD=10) and finally the Pressure Lab had the lowest with (M=22, SD=7).

The quantitative outcome of this process was a set of average ratings by device/environment pair for the 9 collaborative dimensions. Table 5.6 shows the collaboration dimensions and potential correlation with learning gain across the

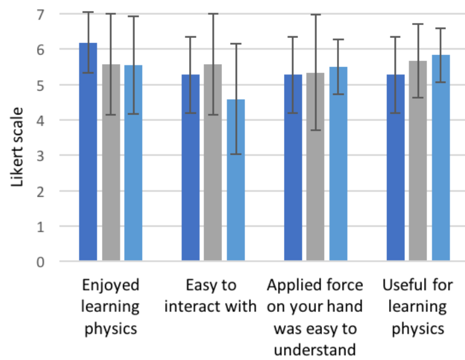

(a) User experience scores

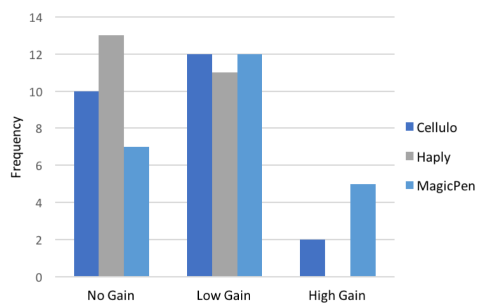

(b) Learning gain v.s. frequency of individuals

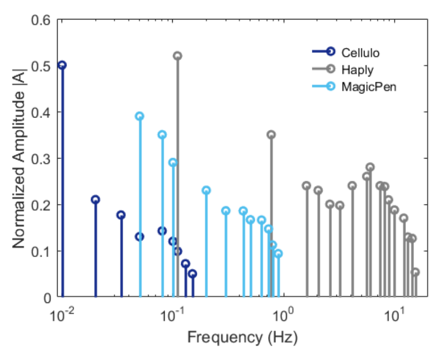

(c) Frequency analysis of force behaviour

Figure 5.7: Summary of quantitative results for the three learning environments of Study 2.

Table 5.6: Correlation between the named dimension (rater's average value for all groups for that device) and the learning gain (measured through pre/post test, of individuals). The dimensions in the first column are taken without modification from [151]. 1,942 cases of collaboration from Study 2 were analyzed; % for each dimension are listed in first column.

Collaboration Dimension	**Example of High Collaboration**	**Learning Gain Pressure (Cellulo)**	**Learning Gain Collision (Haply)**	**Learning Gain Electrostatics (MagicPen)**
Sustaining mutual understanding (SM) (36%)	Confirmation, Asking ideas, Asking for knowledge, Continuing conversation	**Sig** (Partial mediation: Table 5.7)	Non-sig	**Sig**
Dialogue management (DM) (7%)	Smooth flow, Right signaling for having a smooth conversation and turn taking	Non-Sig	**Sig**	**Sig**
Information pooling (IP) (24%)	Presenting new information, Expressing new thoughts	**Sig** (Partial mediation: Table 5.7)	Non-sig	Non-sig
Reaching consensus (RC) (13%)	Explicit agreement/disagreement with rational reasoning	**Sig** (Partial mediation: Table 5.7)	**Sig**	**Sig**
Task division (TD) (3%)	Dividing the task into subtasks, Moving towards the solution step by step	**Sig**	Non-sig	Non-sig
Time management (TM) (0.1%)	Stay focused on the main topic, Evaluate time, Finish irrelevant conversation	Non-sig	Non-sig	Non-sig
Technical coordination (TC) (12%)	Knowledge of tools, Technical skills, Using the tool properly	Non-sig	Non-sig	Non-sig
Reciprocal interaction (RI) (4%)	Respect, Symmetrical relationship, Polite, Invite to converse, Avoid aggression	Non-Sig	Non-sig	**Sig**
Individual task orientation (IT) (1%)	Engagement, Perform the task correctly, Encouraging attitude, Involvement	Non-Sig	Non-sig	**Sig**

three device / learning environment combinations. The examples of high collaboration in Table 5.6 elaborate on how coders scored dyad collaboration. While each of these dimensions can individually contribute to higher learning gain, they can also indirectly impact and support other dimensions.

We continued to search for any correlation between the ranked dimension and the learning gain (measured through pre/post test, of individuals). The dimensional correlation matrices shown by environment in Figure 5.8 suggests varying dynamics of collaboration. These differences could also be caused by different background knowledge, the types of activities and the nature of haptic experiences.

Statistical Mediation Analysis:

Mediation analysis quantifies the extent to which a variable participates in the transmittance of change from a cause to its effect. We argue that participants can learn from (a) collaboration (*e.g.*, learning directly from a partner), (b) haptics directly (by feeling), and/or from (c) collaboration mediated through haptics. As such, we view hCIs as windows into potential mediation as in (c).

Here, we specifically use *mediation* in the sense that we observed instances of certain dimensions during the period of an hCI, and hence infer that their impact is mediated by the hCI. For example, we saw many dyads *Information Pooling (IP)* during an hCI, and subsequently reaching a final answer, thus contributing to a learning gain mediated by the hCI.

Using linear regression we found a significant relationship between the number of hCIs in a session and the learning gains registered by each participant in that session, for two environments (Table 5.7) – namely Cellulo and MagicPen (**Coef**=0.4, **SE**=0.01, **F-Value**= 6.06, **p-value**= 0.033).

Our statistical results show that three of collaboration dimensions in the Pressure environment were *partially mediated* through the use of the haptic device, which means the effect of collaboration still exists on the learning gain. We found no evidence of mediation with MagicPen, which suggests that both collaboration and haptic feedback contributed in learning gain.

Based on direct observation, this outcome is likely explained by the Pressure learning activity inviting greater collaboration than the Electrostatic environment

Pressure Lab

	LG	SM	DM	IP	RC	TD	TM	TC	RI	IT
LG	1.00	0.73**	0.39	0.90**	0.60*	0.60*	0.34	0.32	0.03	0.13
SM	0.73**	1.00	0.17	0.89**	0.84**	0.31	0.56	0.48	0.16	0.12
DM	0.39	0.17	1.00	0.46	0.48	0.50	0.18	0.52	0.47	0.81**
IP	0.90**	0.89**	0.46	1.00	0.80**	0.51	0.57	0.38	0.08	0.32
RC	0.60*	0.84**	0.48	0.80**	1.00	0.23	0.61*	0.62*	0.38	0.45
TD	0.60*	0.31	0.50	0.51	0.23	1.00	0.09	0.46	0.29	0.15
TM	0.34	0.56	0.18	0.57	0.61*	0.09	1.00	0.17	0.12	0.37
TC	0.32	0.48	0.52	0.38	0.62*	0.46	0.17	1.00	0.60	0.20
RI	0.03	0.16	0.47	0.08	0.38	0.29	0.12	0.60	1.00	0.49
IT	0.13	0.12	0.81**	0.32	0.45	0.15	0.37	0.20	0.49	1.00

Collision & Momentum Lab

	LG	SM	DM	IP	RC	TD	TM	TC	RI	IT
LG	1.00	0.27	0.59*	0.21	0.59*	0.09	0.28	0.21	0.33	0.10
SM	0.27	1.00	0.81**	0.47	0.72**	0.75**	0.67*	0.87**	0.85**	0.43
DM	0.59*	0.81**	1.00	0.52	0.89**	0.73**	0.60	0.63*	0.67*	0.60*
IP	0.21	0.47	0.52	1.00	0.51	0.69*	0.38	0.46	0.59*	0.45
RC	0.59*	0.72**	0.89**	0.51	1.00	0.73**	0.81**	0.58	0.60*	0.68**
TD	0.09	0.75**	0.73**	0.69*	0.73**	1.00	0.66**	0.59*	0.55	0.73**
TM	0.28	0.67*	0.60	0.38	0.81**	0.66**	1.00	0.56	0.55	0.57
TC	0.21	0.87**	0.63	0.46	0.58	0.59*	0.56	1.00	0.77**	0.26
RI	0.33	0.85**	0.67*	0.59	0.60*	0.55	0.55	0.77**	1.00	0.27
IT	0.10	0.43	0.60*	0.45	0.68**	0.73**	0.57	0.26	0.27	1.00

Electrostatic Lab

	LG	SM	DM	IP	RC	TD	TM	TC	RI	IT
LG	1.00	0.87**	0.66*	0.44	0.58*	0.29	0.31	0.27	0.62*	0.59*
SM	0.87**	1.00	0.59*	0.54*	0.75**	0.52	0.09	0.27	0.58*	0.52
DM	0.66*	0.59*	1.00	0.47	0.71**	0.67*	0.18	0.56	0.68*	0.52
IP	0.44	0.54*	0.47	1.00	0.66*	0.48	0.18	0.74**	0.39	0.65*
RC	0.58*	0.75**	0.71**	0.66*	1.00	0.63*	0.31	0.57	0.55	0.43
TD	0.29	0.52	0.67*	0.48	0.63*	1.00	0.40	0.61*	0.49	0.33
TM	0.31	0.09	0.18	0.18	0.31	0.40	1.00	0.24	0.09	0.01
TC	0.27	0.27	0.56	0.74**	0.57	0.61*	0.24	1.00	0.37	0.63*
RI	0.62*	0.58*	0.68*	0.39	0.55	0.49	0.09	0.37	1.00	0.44
IT	0.59*	0.52	0.52	0.65*	0.43	0.33	0.01	0.63*	0.44	1.00

Figure 5.8: Correlations among 9 collaboration dimensions as well as learning gain (first row/column), by learning environment/device combination. We computed correlations on the average rating of each dimension across three coders with the addition of dyad's learning gain. Cross-system variation seen in the matrices exposes how collaboration dynamics varied by system.
**. Correlation is significant at the 0.01 level (2-tailed).
*. Correlation is significant at the 0.05 level (2-tailed). We selected a 0.05 threshold for reporting the p value and statistical significance.

Table 5.7: Statistical mediation analysis. The impact of some collaboration dimensions (X) on Learning Gain (LG) appeared to be mediated by Haptic Critical Instances (hCI), specifically Sustaining Mutual understanding, Information Pooling and Reaching Consensus.

Cellulo	**Coef**	**SE**	**F-value**	**p-value**	**Sig.**	**Total (X+HCI → LG)**	**Direct**	**Indirect**
hCI → LG	0.08	0.02	7.32	0.02	Yes	-	-	-
SM → hCI	0.53	0.16	11.43	0.006	Yes	**Coef**=0.53 **5% hCI**=0.29 **95% hCI**=1.12 **p-value**=0.01	**Coef**=0.46 **5% hCI**=0.18 **95% hCI**=1.35 **p-value**=0.12	**Coef**=0.07 **5% hCI**=0.434 **95% hCI**=0.57 **p-value**=0.80
IP → hCI	3.86	1.0	15.0	0.003	Yes	**Coef**=0.55 **5% hCI**=0.36 **95% hCI**=0.73 **p-value**=<001	**Coef**=0.60 **5% hCI**=0.26 **95% hCI**=0.92 **p-value**=0.004	**Coef**=-0.06 **5% hCI**=-0.32 **95% hCI**=0.18 **p-value**=0.50
RC → hCI	6.34	1.62	15.28	0.003	Yes	**Coef**=0.56 **5% hCI**=0.19 **95% hCI**=1.32 **p-value**=0.03	**Coef**=0.24 **5% hCI**=-0.225 **95% hCI**=1.12 **p-value**=0.3	**Coef**=-0.36 **5% hCI**=-0.22 **95% hCI**=0.76 **p-value**=0.14

rather than a factor of relative device affordances.

5.6.5 Thematic Categories

To avoid bias, coders evaluated the collaboration dimensions without knowledge of learning gains. After calculating learning gains (Figure 5.7b), we retrieved our codes to reflect on successful or less effective strategies. At this point we focused on the role of haptics on dimensions more highly correlated with learning gains. Based on these results as well as Study 1 insights, we identified four thematic categories which together capture most of our observations (quotes are from Study 2).

1. Communication and sustaining mutual understanding

In several scenarios, participants asked their peers to perform specific tasks. This was to set up the environment to test a haptic experience, or to confirm a hypothesis. Sometimes, it led to discussion where participants invited opinions or asked for more knowledge.

After the haptic experience, participants would seek to explain the device behavior, *e.g.*, by asking questions or clarifications of their partner. At this point, if their background knowledge was sufficient they could continue the discussion and relate it to their final answer. When it was insufficient, they were unable to proceed and chose wrong answers.

We observed participants were able to communicate their ideas conceptually to their partners, when unable to describe it in technical terms. For instance: *[P21] I know it in my head, but I can't explain physics like on like a basal level, but like, you know what I mean?* This leads us to wonder if there needs to be a bridge to the mathematical definitions of these concepts so that participants can conceptually understand and also technically define such phenomena.

We did not see any evidence to support the idea that haptic experience turns to body-syntonic reasoning as proposed by Papert [174]: a line of reasoning where children imagine themselves to be the manipulatives (*e.g.*, turtle robots) and reason about their own behaviour (how they would move if they were the turtle).

2. Joint information processing and reaching consensus

We observed two levels of reaching consensus. The first was to agree on the perception. The second requires a higher level of understanding and reasoning to answer the activity questions. *[P2] "Not linear." [P1] "That's for sure. So that's probably the first." [P1] "I think it's the second one where the force goes down as you increase this sense." [P2] "Yeah, that's true. That's true."*

However, as participants proceeded, they pass the need to reach consensus on perception level until a point in which one doesn't feel the need to test an activity with the haptic device and reason about the answer based on the previous example.

Participants co-constructed a haptic "measure" to assess their designs, even if it did not completely fulfill the requirements. *[P5] "Cause yours have a slower velocity. So yours works better."* This joint measure helped participants to collaborate to make their design better: *[P5] "Let's just go with yours. You don't want to make it better". [P6] "Yeah". [P5] "The best is going to be the one we are most confident with."*

Participants recalled what they have learned in the past to either predict or explain the behaviour of the Cellulo robots. Few referred to the robots as they tried to use mathematical formula or express scientific terms. *[P2] "in fluid dynamics. We learned that the, the atmospheric pressure and the pressure of all this water it's potential energy and one at the bottom is going to again, be turned into kinetic energy."*

Finally, we observed situations where it was difficult for participants to trust their feelings, complicating consensus. This occurred most often when participants were confident with their background knowledge and found it contradictory to what they perceived. This situation makes it difficult to find a common ground and eventually ends with participants writing individual answers.

3. Interpersonal relationship and engagement

Because it was quick to set up a haptic environment and test an idea, participants became more engaged in the design stage. This engagement and increased communication was reinforced when a partner asked about their ideas, discussing them and conceding to a better design, a process promoted by requiring participation from both learners to operate the device.

Engagement, a sought-after attribute of any learning method, is related to play; our Study 2 analysis specifically found Play as one of the 4 gesture classes we observed (Table 5.4 even though play was not part of the assigned task.

We noticed haptic interactions generating interest and exploration in some participants, which could mean that utilizing haptics for learning could be helpful in making more students interested in the field and more willing to explore ideas outside of the normal syllabus. For instance: *[P13] Oh, so it doesn't work. Yeah. Okay. Let's try the other one. Just for fun!*

We also noticed certain participants came in with a prior lack of interest in physics. Some of them were subsequently not motivated to explore the haptic environment or inclined towards finding optimal solutions. It is possible that using haptics at an early stage to prevent an aversion to the subject could be helpful in subsequent learning experiences, and that use of different representations might not necessarily be helpful in scenarios where participants are already biased against the subject. For example, *[P16] Yeah. It just made me remember in the high school days. I don't want to remember that.*

We speculate that the kinds of engagement and conversation-promotion cues identified above would be particularly beneficial in relationship development for partners who do not previously know one another.

4. Uses of haptics in learning

When learning-focused, participants mainly used the haptic devices for two purposes – learning and revising. Many participants used the device to confirm concepts they were already familiar with, for example: *[[p13] ... , it's more like a review or view. I mean, I think I knew this already. It was just the fact that remembering how it applied sort of like seeing it kind of practically in your life was good refresher.*

However, participants learnt new information as well: *[P20] I learned something new.*

This would be inter-connected to the information polling theme when participants with background knowledge, as those with a thorough understanding of the physics concept behind the experiment would be able to use the haptics environment for confirming their existing knowledge, while participants with no or minimal background in physics would probably learn new information.

Participants used a variety of learning approaches to derive answers as well as communicate explanations to their partner. Some chose to explain certain phenomena in physics using conventional learning methods like physics formulae and reasoning, whereas others explained using new approaches like real-world examples and trial-and-error in the haptics environment. While some could relate the behaviour to the mathematical relationships in the physics formula, we observed that many participants could not remember the exact formula or used a wrong formula to explain some relationships between the parameters which often failed to convince their partner. Real life examples were helpful, *e.g.*, lever or bike gears were more successful in explaining the behaviour of the hydraulic jack in the Pressure environment. We also observed haptic experimentation with quick trial-and-error when participants had some guesses and they were not sure about them.

5.7 Discussion

We can now reflect on what the quantitative and qualitative results of these two studies yield towards our research questions and design guidance. Our analysis – from broad qualitative coding in Study 1 to a focuse on critical haptic instances (hCIs) in Study 2 – went deeply into the what, why and how of haptic use towards collaborative grounding and learning.

5.7.1 Research Questions

RQ1: How can force feedback affect the process of grounding in collaborative learning environments?

Haptic grounding happens: First and foremost, we saw that haptic grounding did occur and was utilized by participants when the context made it available. In Study 1, in the ~30m of analyzed activity per dyad we found an average of 114 codes (for grounding events, haptic gestures and intentions). We broke down these codes per subcategories and report them as percentages. The haptic gestures (Table 5.4) were connected to grounding events and interpreted by intention, with rich examples in the qualitative part of the study.

Platform affordances make a difference: In Study 2, the hCI analysis together with the multiple device/environment combinations gave a qualitative look at what specifically was happening, and how the different device capabilities were supporting those developments.

We detail rich insights into these differences with respect to haptic strategies (RQ2 below) but here identify higher level takeaways about what a specific force-feedback platform can enable:

- Fast realtime response (MagicPen) facilitates shared simultaneous experience.
- Shared rather than person devices can drive increased collaboration, but precludes simultaneous perception).
- A small workspace limits fast, ballistic activities (Haply) which are more possible with large ones, given sufficient update rates (MagicPen).
- Autonomy and ability to stand up is fun and good for many kinds of hypothesis testing (Cellulo).
- Higher fidelity enables perceptual precision and refinement in exploration and experiments; environments need to take device fidelity into account.

Roles and stages: We see evidence for two stage-related roles that haptics enabled for users. Early on, it was helpful in simply mediating their collaboration, as an alternate communication tool that could convey, often nonverbally, ideas or confusions otherwise hard to express. It helped to advance the conversation; and we speculate that for participants who do not know one another well, this could be even more important. At a more advanced stage, and ideally with the benefit of a minimum of scaffolding background knowledge, learners were able to use it as a tool to achieve a learning goal.

RQ2: What haptically enabled strategies do learners use to create mutual understanding?

We got a hint of strategies from Study 1, in terms of intentions behind their haptic gestures (Table 5.5), but greater depth in Study 2's hCI analysis. Participants used haptics for exploration during monitoring and to evaluate and reflect their ideas in order to reach consensus or to resolve disagreement.

Overlapping purposes: Strategies often merged and combined. Study 1 identified co-occurrence of haptic intentions (Table 5.5), with nonuniform linkages and sequencing among gestures carried out with particular intentions (Figure 5.5). Study 2 identified recurring behavior patterns in another way, using 9 collaboration dimensions (Table 5.6) which we examined for co-occurrence through correlation (Figure 5.8), including a preliminary look at trends associated with learning gains. Primary takeaways are that both gestural intentions and collaboration purposes worked in tandem; and that platform impacted collaboration patterns, reinforcing a haptic role in the collaboration dynamics.

Platform and strategy: We also found that the environment/device combinations promoted different strategies. For example, in the Pressure Lab participants coordinated their actions to set-up then observe and compare different (semi-autonomous) behaviours of the Cellulo robots. Conversely, in Collision & Momentum, Haply learners responded more to the perceptual responses to the environment's behavior, *i.e.*, reporting and discussing their perception of the impact when their partner threw a virtual object (using the laptop trackpad) towards the Haply avatar. Such sensations were not forthcoming from the slower-moving Cellulo, nor did MagicPen's structure and activity really invite it. But in the Electrostatic Lab with MagicPen, learners could mutually perceive changes in forces and perceive the effects of their peers manipulation's of the environment in real time, and we saw that this combination of shared perception and fast-realtime responsiveness seemed to enable them to reach consensus more easily than in the other environments.

A similar effect should be available with the Haply, if each collaborator had their own and bandwidth constraints were met by the activity components (they did not permit it in this case). While not the reason we used just one Haply, it did give an opportunity to see what happened when the haptic device had to be shared. The sharing may have amped up their collaboration, since they had to negotiate their use of the single device.

5.7.2 Practical Considerations in Designing to Promote Haptic Grounding

Designers and technologists seeking to employ a haptic channel to support collaboration in dyad work should take encouragement from these results, given their demonstration of participant readiness to employ this channel, their versatile and creative use of it and the evident impact it had on joint interaction strategies. We have tried to shape our findings into some practical, if early, thoughts on things to keep in mind.

Acknowledging and supporting grounding stages in an application's design may empower a modality's use.

We saw active evidence of learners moving between grounding stages – monitoring, diagnosing and repair. Awareness that these stages have different requirements can enrich a design: *e.g.*, providing for private to shared perception; opportunities to confront ambiguities, to trigger discussions and idea-testing; and always, the need for each individual to be able to show or ask about what they or the other is feeling, for comparison with their own.

Activities have a natural rhythm which need to be accommodated

The three stages we developed and followed for our learning activity (Familiarization / Accumulation / Design) highlight a sequence that could be exploited in activity design not only for learning, but for other creative contexts, particularly when merged with grounding stages. While the first and last stage are typically included, we do not always allow enough time on Accumulation - predicting and reasoning for the purpose of understanding, even before we have a design objective. Environments should provide rich opportunities for this step.

To bridge from private reflections to shared understandings, there must be space in the interaction to do both.

Haptics offers a private modality from which to build perceptual bases for grounding;

from there, users can identify and resolve misunderstandings when they occurred at higher conceptual level. It enable users to explore, interact and communicate their experiences.

Haptics can become a voice when verbalization of ideas is a struggle.

Without dealing with semantics and technical terms, haptics offers its own vocabulary to help users achieve mutual understanding. Our observation suggests that activities with haptic environments can come prior to the formal theory classes and help learners to obtain an intuitive understanding of the learning concepts. Lack of technical terms would not burden haptic collaboration between peers.

A little situational knowledge goes a long way; without enough, what one feels may not make sense.

We saw that participants with at least some background knowledge used the device differently, and more purposefully, than those with little or none. There is a threshold of cognitive scaffolding below which sensemaking of the physical sensations is difficult; one cannot accurately interpret behavior and movements. It is likely that a lesson's design can be made adaptable to counter this challenge, using augmenting visuals at the earliest stages to assist novices in forming a usable grasp of the representation. This kind of introductory experience has been suggested for augmented reality in education [184].

To reaching consensus through haptic sharing, start low and don't lose track of it.

We observe two levels of reaching consensus. We realized that agreeing on basic perceptions is important to focus on at the beginning of the activity, when learners are building trust on the device and the information that they receive through it. Maintaining this common ground can build to consensus in increasingly abstract concepts, without the need to return to basic perception.

Haptics can be fun, but it has to be fully accessible.

Haptics can be a fun and stimulating way to engage, learn and interact; we saw it drawing people in. However, it needs to be easy, not frustrating. As for many technologies, the fun/frustrating line will be at a different point for different people.

5.8 Conclusions and Future Work

In this chapter we studied the role of force feedback on grounding, examining haptic strategies that partners employ in collaborative learning contexts. Our two studies sought patterns in participant's haptic interactions, their grounding actions and the intentions behind them, with qualitative and qualitative results revealing how participants naturally use haptics to set-up the haptic environment and test their ideas. We found that participants predominantly employed haptics to explore the environment, communicate their hypotheses and repair possible misunderstandings in an effort to reach consensus. But our analysis also exposed rich, complicated patterns in collaboration dynamics, enabled by environment features and device and device capabilities. These patterns involved intentions such as curiosity, relation-building and fun, and naturally facilitated a non-verbal language. The analysis also exposed the pitfalls of inadequate threshold knowledge, technical frustration, and prior adverse experiences.

In our second study, critical haptic incident analysis provided an objective measure to evaluate different dimensions of collaboration and their correlation with learning gain. Qualitative and quantitative results suggested that use of haptics impacted the collaboration dynamic and strategy differently according to the type of haptic interaction and the learning tasks. When the haptic activity was more collaborative, haptics partially mediated the impact of collaboration on learning. Even when the nature of the task did not required haptic collaboration, we found that haptics and collaboration could separately improve learning.

We reflected on the lessons we learned through our two studies to help haptic designers and educational specialists deliver their haptic information successfully to the users, either to sustain mutual understanding or to create collaborative learning activities.

5.8.1 Limitations and Future Work

From this foundation we anticipate several next steps.

Co-located to remote collaboration

These studies were designed pre-pandemic and executed during lockdown, even as research team members watched young family members struggle with connection-

Table 5.8: Touch in support of collaborative grounding through logical reasoning and factual evidence

Key Haptic cues	**How it is perceived – through an objective lens**	**Example of purpose – information that is sought**
Thermal	Differences between the object temperature and the body temperature	Finding the thermal conductivity in materials [104]
Identifying textures and material surface property	Tactile acuity /roughness/ friction/ finding primitives, and symmetry	Identifying different materials based on their surfaces [156]
Shape and size recognition	Identifying edges/ spatial information and geometries	Students' conceptions of the animal cell [156]
Compliance	Resistance to applied force	Tissue stiffness for surgical robots [115]
Force behaviour	Perceiving changes in force magnitude and direction	Coulomb's law (attractive-repulsive forces)

building in online learning. We are eager to see how our findings translate to situations where remote partners face various reductions in the quality and ease with which they see and hear one another, both faces and what they are doing with their hands. We speculate that the haptic channel grounding benefits that we see in co-located scenarios may be even more valuable in remote ones, and this is an excellent time to find out.

From novelty to skill - longitudinal examination

Few of our participants had any prior experiences with force feedback. Novelty creates engagement, while inexperience obstructs working with the haptics device and interpreting force cues. A longitudinal study can explore how users who become more literate over time can become more effective communicators, explorers and collaborators.

Learning outcomes

Our focus here was to qualitatively understand the role of haptics in grounding, and 24 participants gave more data than we could use. However, more is needed for reliable statistical insights. Ultimately, for our learning use case the prize is impact on learning outcomes, which will require a large sample to determine given the many sources of variation in individual's learning situation.

Other haptic and multimodal cues

Here we only studied the changes in force behaviour; however, many more could be employed in the grounding process. In Table 5.8, we categorize a set of examples to demonstrate the use of haptics in learning activities and how they are perceived through an objective haptic lens. Future studies can investigate grounding with other and combinations of haptics cues, as well as myriad multisensory combinaions with other senses.

Chapter 6

A Framework for Physically Assisted Learning (PAL)

Preface – *Inspired by Harold's purple crayon, in previous chapters we investigate the technical challenges of creating a platform that allows us to investigate the importance of haptics in learning. We proposed a technology (MagicPen) that can haptically render the virtual world similar to Harold's purple crayon. Further, we studied the two core haptic interactions for regenerating Harold's experience of designing and then exploring the hypothetical imaginary world.*

Building on the previous discussions and outcomes, in this chapter we revisit different categories of interaction that are enabled by the force feedback support. We then complete Harold's journey using haptics, by offering a seamless, haptically supported continuum between the activities of constructing an environment (design) and exploring it. We present a Physically Assisted Learning (PAL) framework to achieve a better understanding of how to build a model by drawing it and then exploring the model through the sense of touch. The material in this chapter is taken directly from (Kianzad et al. *2021)*[1].

[1] Soheil Kianzad, Guanxiong Chan, Karon E. MacLean "PAL: A Framework for Physically Assisted Learning through Design and Exploration with a Haptic Robot Buddy," Frontiers in Robotics and AI, pp 228-250, Vol 8, 2021.

6.1 Overview

Robots are an opportunity for interactive and engaging learning activities. In this chapter we consider the premise that haptic force feedback delivered through a held robot can enrich learning of science-related concepts by building physical intuition as learners design experiments and physically explore them to solve problems they have posed. Further, we conjecture that combining this rich feedback with pen-and-paper interactions, *e.g.*, to sketch experiments they want to try, could lead to fluid interactions and benefit focus. However, a number of technical barriers interfere with testing this approach, and making it accessible to learners and their teachers. In this chapter, we propose a framework for Physically Assisted Learning based on stages of experiential learning which can guide designers in developing and evaluating effective technology, and which directs focus on how haptic feedback could assist with *design* and *explore* learning stages. To this end, we demonstrated a possible technical pathway to support the full experience of designing an experiment by drawing a physical system on paper, then interacting with it physically after the system recognizes the sketch, interprets as a model and renders it haptically. Our proposed framework is rooted in theoretical needs and current advances for experiential learning, pen-paper interaction and haptic technology. We further explain how to instantiate the PAL framework using available technologies and discuss a path forward to a larger vision of physically assisted learning.

6.2 Introduction

The learning of topics once delivered in physical formats, like physics and chemistry labs, has moved into digital modalities for reasons from pragmatics (cost, maintenance of setups, accessibility, remote delivery) to pedagogy (topic versatility, personalized learning, expanded parameter space including the physically impossible). Much is thereby gained. However, typically accessed as graphical user interfaces with mouse/keyboard input, these environments have lost physical interactivity: learners must grasp physical concepts in science and math through disembodied abstractions which do little to help develop physical intuition.

Physically interactive robots coupled with an interactive virtual environment (VE) offer an alternative way for students to encounter, explore and collaboratively share and build on knowledge. While contemporary technology and learning theories have not yet delivered a robot system sufficiently versatile to support a wide range of learning needs and environments, we can nevertheless propose and separately evaluate design dimensions that a haptic robot and accompanying interactive VE enables. The objective of this chapter is to facilitate the design and assessment of this new class of learning technology by articulating its requirements via a framework.

Experiential learning theorist – (Kolb ,1984) [122] – posits a four-phase cycle that learners ideally repeat iteratively: concrete experience (CE), reflective observation (RO), abstract conceptualization (AC), and active experimentation (AE). In this chapter we focus on how a haptic robot might be engaged in the stages of this cycle which naturally lend themselves to physical manipulation: **active experimentation**, through *designing* a virtual experimentation environment suitable for a question they have, and **concrete experience**, through *exploring* the environment they configured.

A Vision for Physically Assisted Learning: A Sketch-Based Design-Explore Cycle

The ability to draw a model, then feel it (active experimentation around an idea, then associated concrete experience of it – forming and testing a hypothesis) may be key to elevating interactive sketching to experiential learning. When exploring, learners can extend their understanding of a domain of knowledge by physically interacting with a virtual model – making abstract concepts more accessible, and approachable in new ways. When they are designing, physicalized digital constraints combined with sketch-recognition intelligence can help them to expeditiously express their thoughts by sketching to the system, with the added benefit of representing the resulting model to a co-learner. Finally, exploring one's own designs now becomes a holistic cycle: the learner challenges their knowledge by dynamically posing their own questions and mini-experiments as well as others' by designing models, then reflecting on the outcome of interacting with it.

As a concrete example: to "play with" the dynamics of a physical system (*e.g.*, a mass-spring oscillation), a learner is assisted by a force-feedback-enabled drawing stylus to sketch the system on an arbitrary surface. The system recognizes the drawn ink as, say, a mass connected to a ground through a spring. Using the same stylus, the learner can then "grab" the drawn mass and pull on it. To test a question about parallel versus series spring networks, they can mentally predict then quickly draw the two cases and physically compare the result. Similarly, they could test relative oscillatory frequencies by extending the spring then "releasing" it. By writing in a new spring constant ("K = 2") they can modify the spring constant. The same process can be applied in other domains, such as in designing-to-explore an electronic or fluid circuit, and to improvisationally testing equations defining system properties. This use case (Figure 6.6) and others are implemented and elaborated later in this chapter.

Technical Challenges and Ways Around Them

Aspects of the AE and CE experiential learning stages have been studied and validated in isolation using tangible user interfaces, robots and haptic devices, and the results underscore the general promise of this approach [148, 185, 236]. However, few systems support physicalized interaction in both stages, far less fluid transition between them.

This is at least partially due to the technical difficulties of working with present-day versions of these technologies. For example, conventional grounded force-feedback haptic systems can theoretically support VE creation and interaction, but in practice, they require extensive time and expertise not just to create but even to apply small variants in learning purpose, which often is unavailable in a school setting. Their expense, limited-size and desk-tethered workspaces and single-user nature preclude mobility and collaboration and tend to be too high-cost and require significant technical support. Other robot technologies are mobile and collaboration-friendly, but do not convey physical forces – e.g. a robot puck with which a user can control tokens on a graphical screen.

However, a handheld force-feedback tool that combines a spectrum of autonomy with physical interaction can potentially overcome these technical limitations: *e.g.*, a robotic pen which can assist a learner in navigating concepts of physics and math by conveying physical forces modeled by an environment drawn by its holder. Technically, this system must read and understand the user's sketches and notations, translate them into a VE and associated parameterized physical models, then animate this environment mechanically with a physics engine rendered through a suitable force-feedback display – ideally with the same handheld tool with which they drew the environment. A haptic device in the general form of a handheld, self-propelled and high-bandwidth robot can generate untethered, screen-free experiences that encourage collaboration.

This concept is technically feasible today without any intrinsically high-cost elements, with the haptic pen itself fully demonstrated [113, 116], but significant engineering remains to translate innovations in sketch recognition from other technical domains and integrate them into a full-functioned, low-latency robotic system. Our purpose in this chapter is to consider the potential of this approach based on related technology elements as a proxy for a future integrated system which we know is possible to build if proven worthwhile.

6.2.1 Approach and Contributions

We have designed support based on a theory of activities that has been shown to lead to effective learning, and require this support to meet usability principles suggested by the theory. For example, the cyclical nature of Kolb *et al.*'s [122] learning cycle directs us to minimize cognitive and procedural friction in performing and moving between important cycle activities. Unfettered designing and exploring implies comfortable workspace size and natural command-and-control functions that transfer easily from a student's existing experience – *e.g.*, pen-and-paper diagramming, nomenclature consistent with how they are taught, direct application of parameters, etc. They should not have to switch tools when they switch stages. Meanwhile, their work should be easily visible in a way that teachers and co-learners can see what they are doing and effectively collaborate in their experience [10, 112, 185].

Getting To Confidence that it Could Work

The scope of this chapter is to identify and solve technical obstacles to the instantiation of the theoretically based PAL framework, focusing on the gap in previous work: the connection between physically supported design and explore learning activity, in the form of theoretical rationale and technical proof-of-concept. We need to ensure that the concept's non-trivial realization is feasible, given obstacles ranging from stroke recognition to haptic rendering algorithm and availability of a haptic display with suitable capability and performance.

Only with this evidence will it will be ready to (beyond our present scope) optimize for usability; and thence to evaluate for the pedagogical value of adding

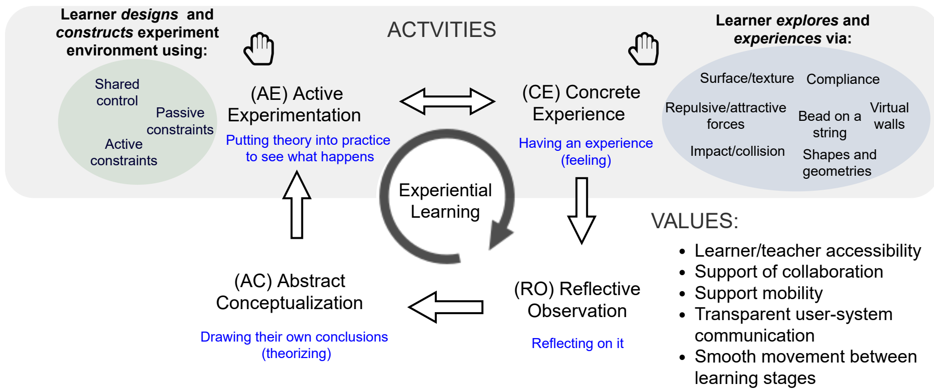

Figure 6.1: The PAL Framework. Physically Assisted Learning interactions haptically adapt stages of experiential learning from Kolb's [122] general framework, with some added features from Honey [93]. Hands-on Active Experimentation and Concrete Experience are most amenable to haptic augmentation, enriching the more purely cognitive Reflective Observation and Abstract Conceptualization.

physical expression and fluidity to the explore-design-explore cycle. Given the complex and individual process of learning, this will require a sequence of user studies to convincingly validate the framework and its impact on learning gain, as well as generalizablity across multiple platforms.

Guiding Support and Assessing Potentials with an Experiential Learning Framework

We propose a **Physically-Assisted Learning (PAL)** framework through which we can systematically compare different candidate technologies' potentials in *unlocking key activities and values* (Figure 6.1). Through the PAL lens, we view learning via the physically supported **activities** of *designing (AE)* and *exploring (CE)*; and assess platforms against key cross-cutting **values** of *learner/teacher accessibility* [171], support of *collaboration*, untethered [114], screen-free *mobility*, *transparent* user-system communication [192], and *seamless transitioning* between learning stages.

We are using PAL as a tool to understand the impact of device attributes on learning strategies and outcomes, as well as collaborative effectiveness, self-efficacy, creativity, and performance in drawing and design. Throughout the chapter, we will relate needs, technical challenges and approaches to this framework, and consider how the candidate technologies stack up on its values under the two activities of focus.

We contribute:

(1) The Physically Assisted Learning (PAL) framework which can (a) conceptually and constructively guide the design of haptic science-learning support; and (b) lead directly to articulation of special requirements for *explore*-type contexts like learning, including fluid access to large ranges of model structure and parameterization.

(2) Demonstrations of (a) means of addressing these needs, for *designing* with innovative application of hand-stroke recognition, and for *exploring* through haptic rendering with a control approach not available in open libraries (namely passivity control); and (b) a technical proof-of-concept system in which *designing* and *exploring* are haptically linked: a user can draw and then feel a virtual environment.

(3) A path forward: An account of technical considerations, challenges and possible approaches to fully realize this paradigm.

6.3 Background

We introduce past work related to the idea of physicalizing digital manipulatives, relevant classes of haptic force feedback technology, challenges in bringing this kind of technology into education environments, and ways in which haptics have been used for related activities of designing and exploring.

6.3.1 Adding Physicality to Digital Manipulatives (DMs) via Robots

Robots are a class of DMs that use motion along with other visual or audio cues to express information. Children can program robots and therefore observe and experience how defining a set of rules results in intentional behaviours in them. This also gives them the freedom to decide what the robot is, based on how the robot behaves. This flexibility potentially helps learners to use the robot as a probe to explore many learning concepts in different contexts [189].

Haptics can empower digital manipulatives by expanding the imagination beyond the motion of a physical robot, in the behaviour of the virtual avatar and respective feeling of force feedback. While users can manipulate the environment, we posit that the visual and haptic cues can reduce the cognitive load of interpreting the abstract concepts and make the haptic digital manipulative more expressive.

Returning to our mass-spring illustration: a physical mass connected to a real spring is a manipulative that can demonstrate the concepts of elasticity, inertia, vibrations and resonance. A programmable robot can visibly implement the mass-spring behaviour through its reactive motion. With physical user interactivity, this robot becomes a ***haptic digital manipulative***. Combined with a graphical display, it could tangibly render the system with learner-specified parameters – shape, size, spring and mass constants – and expose learners to the reaction forces and dynamics of pulling and bouncing it [154] as well as new combinations of springs, and varying viscosity and gravitational force. Such a system can simulate many other physical systems, *e.g.*, gas, fluid or electronic circuits.

6.3.2 Relevant Educational Theory and Design Guidelines

Learning Through Experience

In Piaget's Constructivism [181], knowledge is seen as deriving from individuals' experiences, rather than as a transferable commodity. Learners actively construct and re-construct knowledge by interacting with the world [9]. According to Piaget's cognitive development theory, to know an object means to act on it. Operation as an essence of knowledge requires the learner to modify and transform an object, and also understand the process of transformation; leading to knowledge of how the object is constructed [180]. Several schools of educators [46, 161, 174] have emphasized physicality in educational learning tools and direct manipulation of objects. These theories underlie a goal of providing tools that enable learners to operate on multiple instances of knowledge construction. Papert based his Constructionism on his supervisor's Constructivist [175] learning theory. Constructionism, in addition, takes into account the social and situational aspects of learning. According to Papert, learners will be more involved in learning when they are constructing something tangible (schema) that is shareable and justifiable when other learners can observe and criticize or even use it.

Extending Experience With Reflection

Meanwhile, Cornu [42] propose three iterative steps of *externalization*, *sense-making of meaning*, and *internalization*, through which reflection links experience to learning. Often discussed in social constructionism literature, these steps have been applied to a wide range of human actions in the world and society, including the use of feedback (from people, or the results of physical "experiments") to develop the meaning of the self.

Haptic Digital Manipulatives as Vehicles for Experience and Reflection

The theories above have been applied to a wide range of tangible user interfaces and digital manipulatives. Through educational robots, experiential learning can be tangible and digitally supported, and specifically invite reflection. Resnick's process of *reflection with robots* [190] starts with the construction of a robot-based environment, in which learners make their own "microworld" by programming it, followed by feedback from robots to help them shape and validate their ideas. Such a reflection cycle can be repeated multiple times, deepening the experience [64].

Within early edurobot work, we sought visions for digital manipulatives suitable for more advanced educational topics. We found examples using robots to aid learners in mindful integration or materialization of ideas through the practice of design [5]; and to support exploration of different domains of knowledge or of abstract concepts by making them more accessible or approachable in new ways [172].

Instantiating these principles in a digital manipulative could help them to work as an *object-to-think-with*, wherein learners instantiate their ideas into a physical model through the object, and can debug or extend their thinking model regarding the outcome. The process of analyzing the validity of execution motivates learners to think about their own thinking, developing their metacognitive ability. This results in (a) gaining higher-level thinking skills, (b) generating more representations and metaphors for their understanding, (c) improving social communication and a collaborative atmosphere, and (d) forming deeper understanding of the concept among learners [11, 22].

6.4 A Framework for Physically Assisted Learning (PAL)

The motivation for the PAL framework is to exploit benefits postulated above for a haptic digital manipulative, in learning and in pen-and-paper interaction, and turn them into a versatile and effective digital manipulative. We previously introduced Kolb's four-stage framework for experiential learning [122], on which we have based PAL (Figure 6.1). Here, we lay out PAL's theoretical basis, then elaborate on its components and explain how we expect learners and designers to use it.

6.4.1 Pedagogical Rationale and Components

Learning is iterative: one builds a mental model of a concept by repeatedly interrogating and manipulating a system, forming then testing successive ideas of how it works in a cycle such as Kolb's. Manipulatives are often designed in a way that will support just one part of this cycle – *e.g.*, to create a microworld *or* to directly interact with one. Our premise is that supporting fluid movement *throughout* the experiential learning cycle will facilitate more resilient mental model formation.

Supporting Kolb's Learning Stages with a Haptic Digital Manipluative

Most of the visions in related work, and the idea of robot-supported reflection more broadly, would support at least one out of Kolb's two "acting in the world" phases: Concrete Experience (CE; having an experience) and Active Experimentation (AE; putting a theory into practice). Here, there is an opportunity for intervention, and also for researchers to observe and try to understand what is happening based on the part of the cycle that is visible. The more internal stages of Reflective Observation (RO; reflecting on an experience) and Abstract Conceptualization (AC; theorizing) are crucial, but can be influenced or inferred only through what happens in the other phases, or through post-hoc assessment, *e.g.*, of changes in conceptual understanding.

The PAL framework's mandate is therefore to help educators focus on physical instruments and strategies that will support learners in CE and AE, and eventually to help us insightfully observe them as they do so.

Early works on edurobots have claimed that robots could be beneficial in all four stages. For example, for Reflective Observation (RO), Resnick suggested that through its processing power, the robot could speed the reflection cycle [190] – externalizing/internalizing from hypothesis to result; modifying parameters, conditions and even time. For Abstract Conceptualization (AC), Papert uses gears as an example where learners can use mechanical objects for conceptualizing physics concepts [174].

Kolb argues that the *interaction and manipulation* of tangible objects is an indivisible part of epistemic (knowledge-seeking) exploration, where the learner purposefully changes the learning environment to see its effect and thereby to understand relationships. When suitably framed through availability of multiple perspectives, parameters and factors, manipulation thus might provide at least indirect support for Kolb's Reflecting Observation (RO) stage [9, 57].

However, these claims are as yet unsupported. Limited to findings that have been validated in controlled studies, we conjecture that a DM approach's influence on RO and AC will be indirect.

PAL Components

A useful (that is, versatile) manipulative should be able to provide the basis for productive subsequent reflection and theorizing during both Active Experimentation (AE) and Concrete Experience (CE). Therefore, **we identified *explore* (CE) and *design* (AE) as PAL's key components: activities which a haptic DM must enrich**.

Further support for centering a PAL framework on these two components, as well as clues towards means of implementing them, emerge from other studies of how haptic feedback can support *designing* and *exploring*. Summarizing these, Table 6.1 has two features of particular interest. First, we populated it with just two of Kolb's four learning activities, because we found few examples of attempts to use haptics or other PDMs to directly support reflection or theorizing. Those we did find (*e.g.*, [83, 187, 191]) proposed systems or studies whose results either showed no benefit or were inconclusive.

Secondly, none of the cited studies examined *both* designing and exploring, but treated them as isolated activities. This may have been influenced by the natural affordances of the devices used. For instance, a Haply (in its unmodified state) can be used readily to *Explore*; but to facilitate creation of micro-worlds (*Design*), we felt we needed to hack it – and chose addition of a drawing utensil. In other words, meeting the principles expressed by PAL triggered specific, targeted technology innovation. More is needed to reach the full PAL vision; the framework provides a blueprint to get there.

We believe that PAL framework fits best in Constructionism learning theory. PAL emphasizes on two main aspects of Constructionism. It demands the learner to construct the learning environment, which later they can create an experience by exploring it.

PAL potentially can enhance other schools of learning. From the physiological standpoint, it is possible to find some added values of using the exploration component of PAL for learning by observation. The learner watches how the system does it and then tries to repeat the action. We can also define a reward system so that the haptic feedback attracts and motivates learners towards a certain learning direction with specific objectives (reward-based learning).

We foresee some limitations in using PAL for Cognitivism learning theory. As Cognitivism focuses on information transformation, perhaps the haptic channel is not the most efficient method of communicating information to the learners as opposed to using visual or auditory channels.

6.4.2 Principles for Creating Digital Manipulatives

We assert two overriding principles that guide us in creating versatile digital manipulatives, based on learning theory discussed in Section 6.3.2 as well as observations of learners' interactions both with conventional pen and paper and with haptic/robotic devices, across a range of learning scenarios.

A digital manipulative needs to serve learners in expressing their thoughts (Design)

According to Ackermann [5], "*To design is to give form or expression, to inner feelings and ideas, thus projecting them outwards and making them tangible*". Design enables individual interactions with and through human made artifacts and involves them in the "world-making" process [75]. The purpose of design goes beyond representing just what exists, by bringing imagination into this existence

Table 6.1: Summary of research informing the use and benefits of haptics in learning, organized by the PAL framework's two activity components. [+] indicates a positive benefit, or [-] no added value was found.

Haptic Benefits	*Design* (Active Exploration)	*Explore* (Concrete Experience)
Understanding and manipulating geometry	[+] [231] Drawing accurate geometric shapes. [+] [166] Computer assist collaborative drawing of different shapes. [+] [140] Increasing the passive stylus affordance through haptic guidance.	[+] [172] Identifying different shapes and number of edges. [+] [156] Understanding the structure and function of the cell membrane transform [+] [105] Learning morphology and dimensionality of viruses; diagnose mysterious viruses by pushing, cutting and poking.
Improving accuracy and speed	[+] [116] Improving accuracy of drawing objects through force feedback assistance [+] [222] Using haptic feedback in a calligraphy simulation reduces writing errors and improves writing speed. [+] [231] Drawing accurate geometric shapes.	[+] [165] Enhancing completion time and interactivity of bimanual tasks. [-] [55] Users were unable to sculpt forms to produce acceptable curved surfaces using haptic feedback. [-] [19] Haptic human–human interaction does not improve individual visuomotor adaptation.
Engagement	[+] [84] Significant increase in students' engagement during the learning activity. [+] [233] Increasing engagement in word-writing activities. [+] [127] Increasing confidence and achieving more realistic drawings.	[+] [220] Enhancing interactions with objects in Augmented Reality [+] [217] Providing realistic sensation of physical interaction in a virtual environment [+] [112] More engagement in educational robotic activities.
Accessibility (*e.g.*, in face of disability)	[+] [164]) Re-learning to write after a stroke. [+] [100]) Haptics improves task performance of children with physical disabilities (review paper.	[+] [221] Allowing visually impaired users to perceive data with greater speed and efficiency.
Understanding of underlying concepts	[+] [142] Designing an optimum system/model by receiving on-the-go force feedback.	[+] [148] Conceptualizing electrostatic concepts through the sense of touch. [+] [235] Building electrical circuits with one or two bulbs. [-] [188] Haptics did not add to learners' ability to understand pendulum principles. [+] [236] Understanding mass-beam balance.

[5].

For example, we often use pen and paper to write down fast-travelling ideas in our minds. Our immediate drawings can reflect our thoughts, experiences and emotions. Particularly for children, drawings reveal the hidden transcripts of their interpretation of the world. From scribbles to detailed, elaborated productions, sketching is both intellectual play and can help us form, develop and communicate our thoughts, a key part of a conceptual process. Sketching is direct, improvisational, expressive, resists distraction, and may promote deeper cognitive processing. Projecting our ideas onto paper makes our thoughts more tangible, shareable, and justifiable; This enhances our communications with others. A versatile manipulative should work as a medium to exchange information between a user and a computer interactively.

These prior findings and observations support the premise that aid from a suitably configured and supported physical digital manipulative can directly impact the active experimentation phase: specifically, when learners are hypothesizing and planning small tests. The environment altogether should encourage the learner to hypothesize, construct a experimental micro-world and set the conditions for the environment, anticipate the result and test it; and iterate to improve their hypothesis.

A digital manipulative needs to support exploration of domains of knowledge (Explore)

Two classes of manipulative proposed by [190] include *Frobel* Manipulatives (FiMs) to model the world, *i.e.*, provide an intuitive way to experience many concepts in physics by making them more accessible (wooden sphere and cube to feel the natural differences between shapes), and *Montessori* manipulative (MiMs) to model abstract structure – *e.g.*, form an approachable way to make math, and geometry concepts more tangible (golden bead materials used for representing number). Haptics researchers show that even a 1D haptic device can support both of these classes when it works as haptic mirror [154], to mimic physical experience, or as a haptic bridge, connecting a dynamic visualization of a mathematical concept with a haptic representation [44]. A versatile manipulative should support both classes using physical interaction with the virtual world through force feedback.

Perhaps the most studied aspect of digital and physical manipulative is the role of physicality in simulation learning for concrete experience (CE) stage. Here, learners try out the action and have a new experience. Through physicality, learners can obtain more embodied experiences and perceive information through touch.

6.4.3 Using the PAL Framework

Learner's Use

Some examples illustrate PAL's two conceptual activities, wherein a learner constructs a microworld then explores it.

Design: The learner must be able to fluidly express rich information to the system. Assistive force feedback to users' pens while sketching can help them manifest and communicate their ideas to other people and to a computer: it might be more efficient and natural if they can feel virtual constraints that support them in generating smooth curves and straight lines as they draw – on a computer screen, paper, whiteboard or other surface. In the future, we can exploit this design space to empower learners to actively design, make, and change their learning environment based on their hypothesis.

Explore: The tool must provide rich sensory information to the learner. The addition of haptics to a digital manipulative (beyond motion alone) potentially supports a more compelling interpretation so that learners can predict and reason about outcomes based on what they feel as well as see.

In this project we explore these two PAL activities – requisite attributes for an object to think with – along with the connection between them. Although such a device could also be seen as an object to promote computational thinking [99] we saw it differently. A DM exploits the computational power of the computer to speed up the *learner's* reflection cycle, which leads to more constructive failures [36]. Throughout this process, learners can explore a variety of representations and solution methods. If followed by a consolidation and knowledge assembly stage, together they can create a productive failure process [108].

Education Technology Designer's Use

Ideation of Form and Prediction of Haptic Value

Designing technology solutions for learning requires ideating innovative concepts and ideas, but also evaluating and prioritizing them. PAL can help inspire educational technology designers with new ideas, and to understand the potential of adding haptics to a particular domain or context. In addition, our implementation shows a technical example of how to use emerging technological capabilities to solve particular problems.

Setting Requirements and Evaluating the Result

PAL can help designers identify *requirements* via experiences that their technology needs to support. Based on Figure 6.1, a designer can create an opportunity map by examining connections between the stages of learning and activity type. For example, to support collaboration in learning electrostatic forces, a learner can construct the environment (design) by placing the point charges; then invite their partner to experience them (explore). A designer can then focus on finding the haptic controls and feedback which will allow the learner to place the point charges correct places (*e.g.*, equidistant), and how to render the force behaviour as learners move respectively to each other.

Based on these requirements, in *evaluation* a ed-tech designer simply needs (at a first pass) to verify that the requirements are being met when learners interact with the system. Are they able to construct the environment, and then place the charges correct? Can a partner experience this? Is the whole experience engaging and usable enough to invite this kind of collaboration? With the assurance provided by intermediate goal and usability evaluation derived from theory-based guidelines, they will be in a better position to proceed to assess how such a system is influencing learning outcome.

6.4.4 First Step: Need for a Technical Proof-of-Concept

In past research supporting haptic *design* and *explore* activities (Table 6.1), what is missing is the *connection between* them. This requires a technical means by which to understand the user's imagination and dialogue in *design* and then bring it into existence by defining its physical, haptic behaviour for *exploring*. For example, if a user draws a microworld consisting of a set of point charges, we need to define the force behaviour of the point charge and make it interactive so that users can feel the forces as they move in the environment.

Once such a system exists, it can misfire for purely technical reasons. For example, expanding the user's available possibilities during *design* – *e.g.*, allowing them to cover a greater variety of concepts in more ways – often introduces new issues such as triggering vibrational instabilities which naturally accompany haptic rendering of dynamic environments with large uncertainties.

In summary, the challenges here are to (1) make an intelligent system that can take unconstrained drawing as an input, and (2) robustly render a wide range of haptic environments with high quality. For the first, advances in artificial intelligence go far in allowing us to infer and display interpretations of user's drawings [21, 47]. For the second, the field of haptic rendering can contribute advanced control methods which when carefully applied should be able to describe and within bounds, to address the environments that may arise when a user is permitted to create ad hoc environments [49, 82].

Putting these elements together is, however, a substantial systems-type contribution, and its initial appropriate validation is in technical performance assessment with respect to force and stability outputs relative to known human psychometric capabilities rather than a user study of either usability or learning efficacy. In the following, we will describe and assess performance of our technical-proof-of-concept system which implements this missing, connective aspect of our proposed PAL framework.

6.5 Haptically Linking the Expression and Exploration of an Idea

Currently available processes for generating and modifying content for haptic interaction impose logistic and cognitive friction between ideation in the form of sketching a problem, idea or experiment the learner would like to understand, and testing that idea in the form of a physicalized model. We aim to reduce this friction.

After describing the technical setup we will use to demonstrate our ideas, we will work through a series of technical instantiations which support increasingly powerful and wide-ranging cases. Each begins with an education use case illustrating how this level of haptics could be useful. Readers may find the first (rendering a haptic wall) distant from our final goal; we have included it as a means of gradually exposing layers of haptic technology needed to understand more complex implementations. While all of the haptic rendering algorithms described here are well known, we show how they can be combined in new ways with other technical features (*e.g.*, stroke recognition) to meet technical challenges that arise from the requirements of a versatile, unrestricted learning environment.

6.5.1 Technical Proof-of-Concept Platform: Haply Display and Digital-Pen Stroke Capture

The demonstrations described here use the Haply Robotics' pantograph system (Figure 6.2, https://haply.co/, [65]) and its hAPI software library [66]. The Haply is a low-cost pantograph, relying on 3D-printed parts which together with good-quality motors and fast communication can offer convincing haptic rendering with respect to accuracy, force levels, responsiveness and uniformity across its 14x10cm workspace (https://haptipedia.org/?device=Haply2DOF2016). It communicates sensor and actuator data via USB to a VE running on a host computer, typically using the Processing computer language. The hAPI library renders haptic interactions by reading the pantograph's end-effector position (moved by the user's hand) and computing output forces sent to two driver motors.

To capture users' sketch strokes, we used a watermarked paper and a digital pen (Neo Smart Pen, [98]) connected to the Haply end-effector. The digital pen captures detailed information about the user's stroke: absolute position, pressure, twist, tilt and yaw. The Neo pen requires watermarked paper, creatable with a standard laserprinter by printing encoded dot files. For erasability and re-usability of sheets, we laminated the watermarked paper and positioned it under the Haply workspace. We calibrated the digital pen's position data with the Haply's encoders. With this system, the user can draw on the laminated paper and the strokes are captured, sent to the host computer and imported to the Processing application that interacts with the Haply.

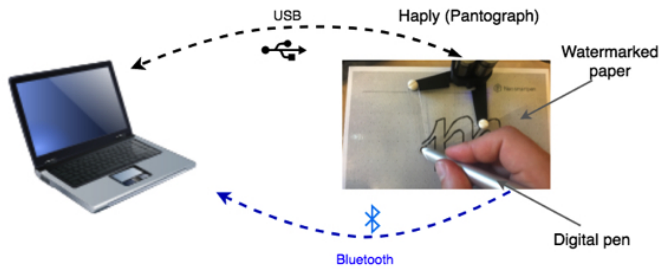

Figure 6.2: Technical Setup. Our demonstration platform consists of a Haply force-feedback pantograph, a USB-connected digital pen, and a host computer. The Haply communicates position information to the host computer and receives motor commands through a USB port. A digital pen captures and conveys thes user's stroke, along with data opressure, twist, tilt and yaw.

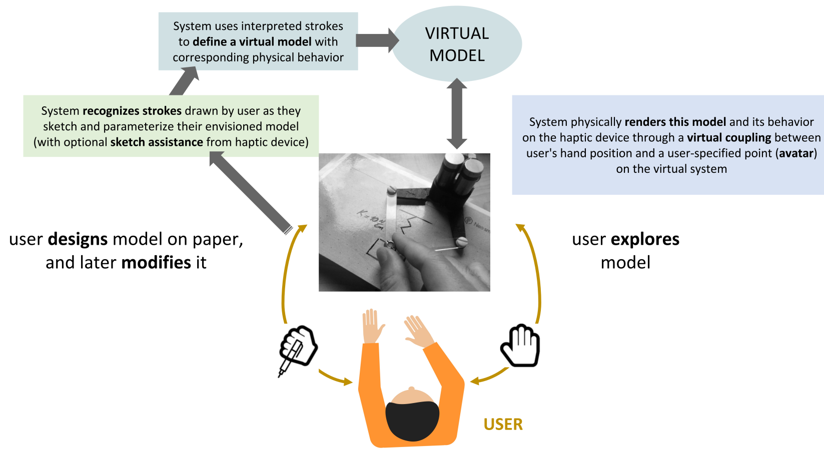

Figure 6.3: Technical implementation required to support *Design* (green) and *Explore* (blue) learning activities in response to ongoing user input. Details are explained in Section 6.5.3.(The user's graphic from Can Stock Photo, with permission).

6.5.2 Level 1: Rendering Rigid Surfaces and Tunnels

We begin by illustrating how haptics could potentially support learning (in motor coordination) with basic haptic rendering techniques.

Use Case: Handwriting Training Guided by Virtual Walls

Past research on motor training, *e.g.*, post-injury rehabilitation, has elucidated effective strategies for utilizing physical guidance, whether from a human trainer or a programmed haptic appliance. Full guidance of a desired movement does not typically produce good transfer to the unguided case; some studies suggest better results by physically obstructing the desired movement in the face of visual feedback, causing exaggerated motor unit recruitment [94, 143]. Learning and improving handwriting similarly involves training numerous haptic sensorimotor activities; these employ both fine (fingers) and larger (arms) motor units. It entails significant mental and motor-control practice, particularly for individuals working against special challenges, such as dysgraphia which can impact 25% of the school-aged population [79, 206].

However, learning and improving handwriting is also a cognitive practice, and often practiced by the young where engagement is also important. Rather than learners comparing their results to a standardized specified outcome, an expert may be able to conceive of better individualized support (more specific, or advanced at a different rate) but requires a means to convey it to the learner as they practice on their own [68].

The priority may thus be easing an expert's customization of exercises, to support repeated self-managed practice [79]. The expert might want to modify details of visual cue presentation and the level and form of haptic guidance [125, 215]; or temporally adapt by reducing force feedback aid over time through control-sharing [116]. Effective feedback must convey correct movements, notify a learner when something goes wrong, and show them how to correct their movement [10]; [8]. Haptic guidance could potentially provide these needed cues when the teacher is not present, without demanding a high cognitive load.

In the PAL framework, the teacher would use the *design* stage, then *explore* to ensure the force feedback works correctly. The learner would access this resource in the *explore* stage.

Here, we show in a basic example targeting elementary school students how a teacher can define a channel within which the learner needs to stay as they trace a letter. This channel will be rendered as a pair of enclosing and guiding haptic walls. This simple demonstration does not attempt best practices for handwriting training, or demonstrate many customization possibilities; it primarily introduces a important building block of haptic rendering, but is also a placeholder for the advanced ways listed above that haptic feedback could be used to customize handwriting support.

Defining a Wall

There are many ways to define a boundary to a computer program. We require a means that is convenient for a teacher or therapist. Working in a context of pen-and-paper, we let the teacher sketch the path which they wish the learner to follow. Their strokes are captured as a time-based set of point coordinates. These can be used either directly, if the stroke sample density is adequate, or with a smoothed line fit to them. We collect the user's strokes as a two-dimensional array, then re-sample it with spatial uniformity and present the result as a one-sided wall. A user can move freely on one side of the wall; if they penetrate the wall from the free direction, they will feel resistance. A teacher can draw a set of one-sided walls as a letter-shaped tunnel to guide a learner in their handwriting practice.

Feeling the Wall: Virtual Coupling

The simplest way to haptically render a wall is to sense the position of the user or haptic device handle, hereafter X_{user}, and compare it with the wall boundary X_{wall}. If X_{user} has penetrated X_{wall}, the penetration distance is multiplied by a stiffness K defining the force that pushes the user out of the wall (Figure 6.4(A, upper)). However, we typically want to render stiff walls, while limitations of haptic device force output and sampling rate create a result which is both squishy and unstable [71]. As shown in Figure 6.4(B-C), increasing K makes a more rigid wall but at the cost of unstable oscillations.

Virtual Coupling (VC) for Stiff Yet Stable Walls

An accepted technique for stably rendering stiff walls, *virtual coupling* connects the haptic end-effector position X_{user} to a point representing it in the virtual world which we define as its **avatar** (X_{avatar}, [197]). A VC links X_{user} to X_{avatar} through a virtual damped-spring, as shown in Figure 6.4(A, lower). A stiff VC spring connects the operator more tightly to the virtual model; they can feel more detail, but it can lead to instabilities.

Thus, a VC's parameters (stiffness and damping) need to be tuned to model properties, such as virtual mass and spring magnitudes, device force limits and anticipated interaction velocities. When these are known and constrained to a limited range, a VC can work well. The VC implementation in the hAPI interface library enables users to change VC parameters [66].

A virtual coupling is closely related to a proportional-derivative (PD) controller, perhaps the most basic form of automatic control structure. The key goals in tuning either system are to (a) set damping to the minimum needed for stability, to limit energy dissipation and consequently responsiveness; balanced with (b) sufficient stiffness to achieve satisfactorily tight connection to the user's motion. System stability is also challenged when the mass of the virtual entity to which the avatar is either bound or touching is too small, or when the system's update (sampling) rate is slow compared to the dynamics of the system (either the virtual system or the user's movement) [204].

Wall Performance in Letter-Drawing Use Case

In Figure 6.5, we show the various mechanisms by which a teacher can define and revise a shape which they want a learner to trace (A-D). In (E), we show an example of a learner *exploring* the tunnel defined by the letter outline, including the haptic rendering performance of the virtual coupling as a learner practice to write an *m*. The spring-damper VC filters high frequency force variations and creates smooth guidance as the user slides between and long the walls; the forces keep them within the tunnel. The user's actual position sometimes goes outside the wall, but their avatar remains within it and the learner feels restoring forces pulling them back inside. Depending on velocity, the user position and avatar may be slightly displaced even while within the wall, as the user "pulls" the avatar along through the damped-spring coupling.

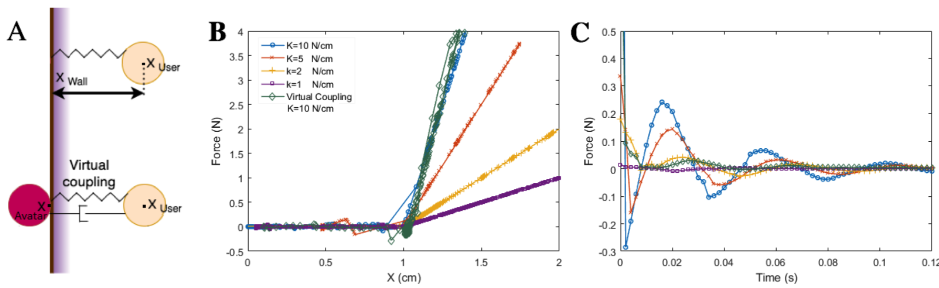

Figure 6.4: Rendering a haptic wall, using a virtual coupling to achieve both high stiffness and stability. **(A) Algorithm schematic**. (*Upper*) In the simplest rendering method, force depends directly on the distance between the virtual wall and the user's hand (haptic device) as it penetrates the wall: $F = K(X_{wall} - X_{user})$. (*Lower*) A *virtual coupling* establishes an *avatar* where X_{user} would be if we could render a wall of infinite stiffness, and imposes a virtual damped-spring connection between X_{user} and X_{avatar}. **(B) Force-displacement behaviour** when the wall is rendered as a direct stiffness or through a virtual coupling. The VC used here also uses the maximum $K = 10N/cm$, and achieves a similar stiffness as when this K value is used on its own. **(C) Oscillatory behavior** of the conditions from (B). In direct rendering, instability increases with K, but with a VC, a high K is as stable as the softest direct-rendered wall. (B) and (C) show data sampled from a Haply device.

To extend this example, a teacher could adjust the tunnel width (a step amenable to parameterization) to customize the experience for the learner. The activity can optionally be visualized graphically, or be done entirely on paper. Learner progress can be quantified through statistics of position error (distance between the physical and virtual avatars) and the force magnitude generated in response to this error.

6.5.3 Level 2: Drawing and Feeling Dynamic Systems

Our second example implements more challenging stroke recognition, and addresses the situation where a virtual coupling is inadequate because of the range of properties that the user may need to access in their design and exploration.

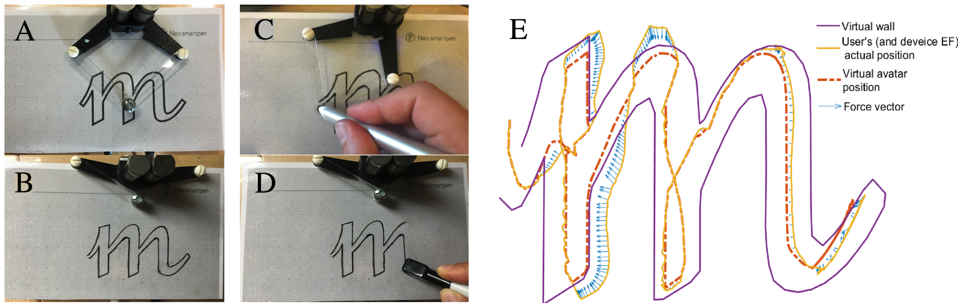

Figure 6.5: A teacher prepares a handwriting activity by defining a letter shape *m*; the learner will then attempt to form the letter with assisting guidance. To create the *m*, the teacher can **(A)** laser-print a computer-generated graphic on paper, **(B)** draw it by hand, or **(C)** manually draw it with haptic assistance. For erasable media, *e.g.*, pencil on paper or marker on whiteboard, the teacher can **(D)** erase and draw a new exercise. **(E)** Exploring the *m* with ink marks rendered as virtual walls.

Use Case: a Mass-Spring System

Hooke's Law is a linchpin topic in high school physics: along with gravity and friction, students learn about the relation between applied force and the amount of displacement in springs and other stretchable materials. They further must be able to define what a spring constant is, how to compute a net constant assembled through parallel and serial spring assemblies, and with support from their teacher, conduct experiments to verify spring-stiffness hypotheses [70]. Here, we use a dynamic system consisting of coupled mass and springs to demonstrate the construction of and interaction with a physical system model based on the PAL framework (Figure 6.6).

System Interprets the User's Stroke

We used a 2D recognition library implemented in Processing (the *$ 1 Unistroke Recognizer* [230] to translate user sketches into a virtual model. *$ 1* is an instance-based nearest-neighbor classifier with a 2-D Euclidean distance function. It can accurately identify 16 simple gesture types, *e.g.*, zigzag, circle, rectangle. To improve performance and customize it to shapes relevant to models our system supports, we created a database to which learners can add their own labeled strokes.

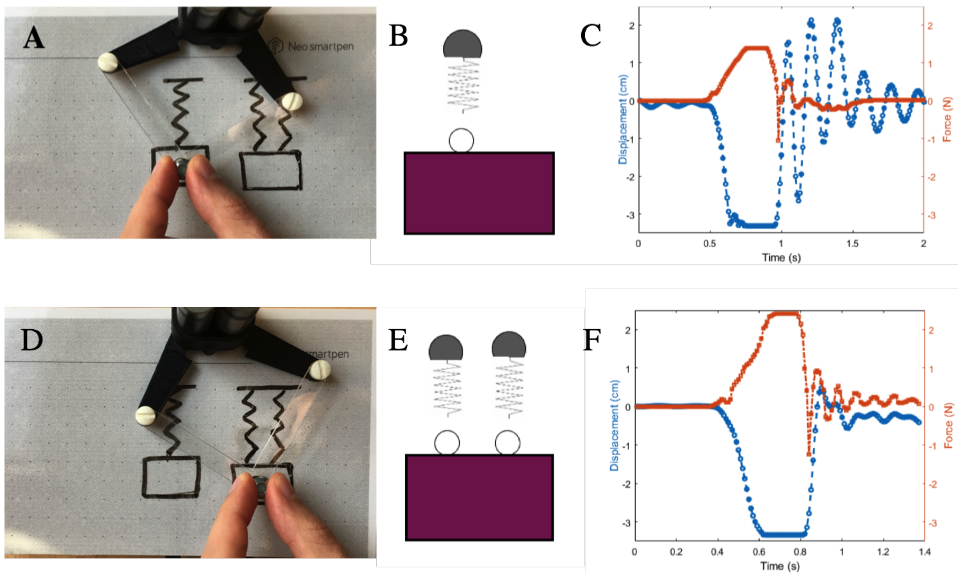

Figure 6.6: Use case: comparing the dynamic behavior of different spring–mass system configurations by drawing then feeling. **(A)** The user sketches a pair of spring–mass systems using a system-readable notation. **(B, E)** Our system recognizes the user's strokes and incorporates them into virtual models. The user can now "connect" to one of the drawn masses by moving over it and *e.g.*, clicking a user interface button. **(C)** Behavior when connected to the single-spring configuration (A). The system implements the corresponding model (B) by pinning X_{avatar} to that mass. The user can then feel the oscillatory force behaviour by "pulling the mass down," extending and releasing the spring. **(D)** The user connects to the two-parallel-springs configuration, and compares its behavior (model E) to the first one. **(F)** compared to **(C)** shows a higher force for the same displacement, and a different oscillatory behavior. This system is implemented using a passivity controller to allow a wide range of M and K values, which are modifiable by hand-writing new values on the sketch.

In the current implementation, the system starts in a training mode where users draw then type to label their sample; then exit training mode and start designing their experiment.

Our current implementation is modal: it needs to know what kind of a system a user is sketching in order to recognize their marks. A zig-zag could represent a spring in a mechanical system, or a resistor in an electrical circuit. This can be done by manually writing the system type's name on the paper with the digital pen as shown by [142] – *e.g.*, , "Hydraulic lab" triggers a hydraulic simulation. The Tesseract optical character recognizer (OCR) system is one of many robust solutions [111]. For simplicity, we selected environments using a graphical user interface.

Reliance on a set notation for sketching has a potential as usability feature or pitfall. If the notation is well known (*e.g.*, taught in the curriculum), it gives the learner a pre-existing language; versus unfamiliar, unmemorable or uncued (*e.g.*, no "tool-tips"). We did not focus on usability refinement at this stage; ensuring it will be an important future step.

System Interprets User Strokes for Model Construction and Parameter Assignment

Ease of environment specification and modification is an important PAL principle. One way that users can specify environment parameters is in the way they draw them. For a mechanical system, a box indicates a mass; mass magnitude is interpreted as the area within the box. Spring stiffness is assigned based on the zigzag's aspect ratio. Haptics can provide assistive guidance to create more accurate drawings. Here, haptic constraints help the user follow implicit geometrical relationships such as relative locations and sizes, through "snapping"; thus the user can perceive when they reach and move beyond the width or length of the previously drawn spring.

Some parameters are harder to indicate graphically, or the user may want to modify an initial value. This could be handled by writing an equation: *e.g.*, set the value of gravitational force with $g = 9.8m/s^2$, or change a spring constant by $K_1 = 10N/cm$. As before, recognition can be done with an OCR like Tesseract, a possibility already demonstrated by at least one other system [142].

Unconstrained Experimentation Requires Stepping Up the Control Law

Fluid exploration means that a learner should be able to observe and feel an object's behaviour and reason about it. This requires changing object properties, comparing behaviour between versions of a model and reflecting on the differences.

Above, we introduced the concept of an avatar as key to rendering a wall through a virtual coupling. The avatar's existence was transparent to the user, its existence implicit in their movement. But when we advance to interacting with multiple dynamic systems – to compare them – users must get more explicit with their avatar. To "hold on" to and interact with a part of a virtual model, such as a tool or to probe part of a dynamic system, they must hitch or pin their avatar to that model element, just as they might when selecting a character in a virtual game.

The combined functionality of (a) pinning and unpinning one's avatar to arbitrary system elements, and (b) allowing unconstrained parameter assignment, is a major departure from how a model intended for haptic rendering is typically constructed. Normally, we design an environment with particular components, set its parameters to a pre-known range, and expect the user to interact with it in a particular set of ways – always connecting through a particular avatar linkage. For example, in a surgical simulation, we might have a defined set of tools, and known tissue parameters. Bone and liver have different properties, and rendering them might be highly complex and computationally expensive, but their properties are known in advance. We can tune a controller (such as a VC) to work with those constrained conditions.

This is no longer the case if parameters can be changed arbitrarily and on the fly, and as usual, the result will be instability. Commonly, several factors can cause instability, such as quantization, delays, and virtual object properties like stiffness and mass. We address this next with the passivity controller.

6.5.4 Level 3: Expanding the Range of Parameter Exploration through Passivity Control

To move beyond the simple tuning heuristics above, we reference the notion of *passivity*. A real-world, nonvirtual system like a wood tabletop or mechanical button or doorknob is *energetically passive* – it will never vibrate unstably when we touch, tap or wiggle it because such oscillations require additional energy which they cannot access. The only energy flowing into the interaction comes from our own hand. At best, we can excite a mechanical resonance (*e.g.*, by bouncing a rubber ball, or pumping our legs on a swingset), but this cannot grow in an unlimited way because of the lack of an external energy source.

In contrast, a haptic display is *energetically active:* it accesses an external energy source through its controller. This is necessary for the system to physically simulate a VE's dynamics. However, instability – often manifested as vibrations that grow without bounds, or unnaturally 'buzz' upon operator contact – occur when the total energy entering the system from the human operator and the controller's commands is greater than the energy leaving it.

Passivity theory underlies a type of controller which can be designed so as to guarantee stability in systems interacting with humans [38, 39, 153]. In essence, passivity controllers bound system movements based on the energy flow through the system: they guarantee overall system passivity by ensuring that the energy input exceeds outputs. It also can achieve global stability through local passivity in subsystems separately. As a result, if we know that other parts of the virtual model and physical device are operating in a passive range, we can focus on the subsystem that the (less predictable) user is interacting with.

Passivity Controller Overview and Design

We designed our *passivity controller* (PCr) with the method described by [86]. Our contribution was mainly to implement the PCr on MagicPen and to evaluate its performance. We made no changes to its design. In overview (Figure 6.7), the PCr is interposed in series between the haptic interface and VE. This location is similar to the virtual coupling controller, and like the VC, the PCr works by acting as a virtual dissipative element; the PCr differs from a VC through its more targeted energetic accounting system.

The human operator interacts physically with the haptic device in continuous time; however, since the control system is digitally sampled, the VE is computed with a time delay typically specified at 1/10 of the fastest dynamics in the system. The human operator is conceptualized as an *admittance* – a source of *flows* (*i.e.*,

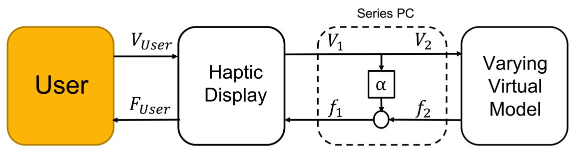

Figure 6.7: Simulation model of a complete haptic interface system and passivity controller, as implemented here. (Reproduced from [86], Figure 8. System blocks are (left to right): user, haptic display, passivity controller α, and virtual environment.

movement), and sink of *efforts* (*i.e.*, forces) – and the VE as an *impedance* – a source of efforts and sink of flows.

At the heart of the passivity controller is α, which is in turn based on the *Passivity Observer (PO)*. The PO, also known as the parameter E_{obsv}, computes the total energy observed in the system at a given moment *(n)* is the summation of energy from the initial time until the moment *(n)*, and can be expressed as:

$$E_{obsv}(n) = \Delta T \sum_{k=0}^{n} f(k)v(k) \tag{6.1}$$

where ΔT is the sampling time, (f) and (v) are effort and flow (force and velocity) of the 1 port network at time step n. Specifically, f_1 and v_1 are effort and flow for the haptic display, while f_2 and v_2 are for the force computed from the VE computation.

When $E_{obsv}(n)$ is positive, the system in losing energy; for negative values it is generating energy. The role of the passivity controller is to create a dissipative element based on the energy generated by the system (The mathematical proofs are presented by Hannaford *et al* [86]). We compute α as:

$$\alpha(n) = \begin{cases} -E_{obsv}(n)/\Delta T v_2(n)^2, & \text{if } E_{obsv} < 0. \\ 0, & \text{otherwise.} \end{cases} \tag{6.2}$$

After the VE model is updated, its subsystem forces are recalculated, then passed through α before being passed as commands to the haptic display's actuators. f_1, the haptic display command force, is computed as the VE force plus the passivity control component (which is acting to siphon excess energy out of the system). For each timestamp (n) we have:

$$f_1(n) = f_2(n) + \alpha(n)v_2(n) \tag{6.3}$$

$$v_1(n) = v_2(n) \tag{6.4}$$

In this implementation

- If the amount of force exceeds the motor force saturation, we subtract the excess amount and add it to the next time step,
- If the user spends significant time in a mode where the PCr is active (dissipating considerable energy to maintain stability), energy will accumulate and the PCr will not transmit actuation forces until the user has backed away from the dissipation-requiring usage, allowing the PCr to discharge. In practice, we reset the PO's energy accumulation to zero every 5 seconds, scenario-tunable scenario or adapted automatically.

Passivity Controller Performance

Example 1: Large-load Coupling

In our first assessment, we examine the performance of our passivity controller for a simple scenario in which the user's position (X_{avatar}) is "pinned" to a virtual mass as if holding it in their hand. We evaluate performance with two load levels and show how the PCr performs on a large-load coupling.

Virtual Coupling: (M=1X) Figure 6.5.4(A) shows the displacement (upper) and energy output (lower) of the virtual coupling system of Section 6.5.2, *i.e.*, without the PCr. The VC parameters are optimized for this system. Thus, when the user (X_{user}) moves 2cm, X_{avatar} follows smoothly with no overshoot, achieving steady-state by 150ms. The maximum kinetic energy of PCr can potentially reduce performance in a normal case where it is not needed, as it may siphon off system energy even when not necessary, being a conservative approach. Therefore for cases close to the system parameters for which the VC was originally tuned, we switch it off.

Virtual Coupling: (M=20X) To understand the effect of changing the virtual avatar properties, we investigate a scenario of increasing the mass of the virtual free body

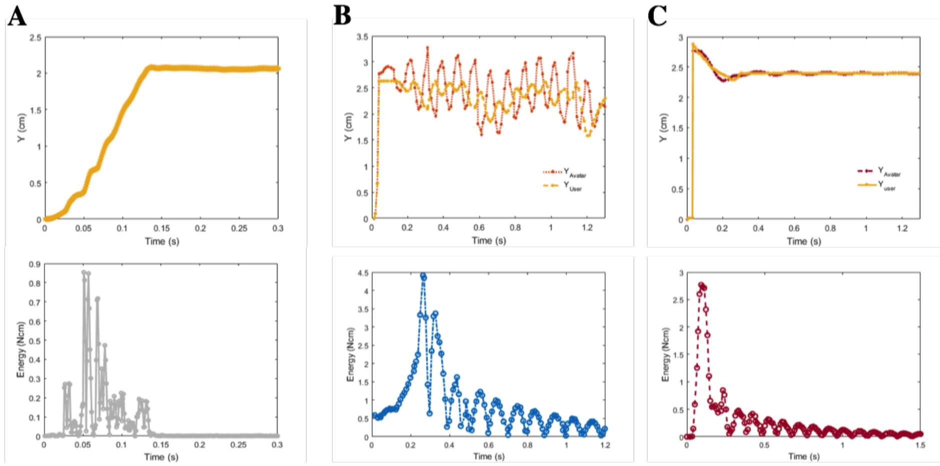

Figure 6.8: Abrupt movements of varying loads. The position, *i.e.*, X_{avatar} as it tracks X_{user}(upper) and kinetic energy (lower) of the load for **(A)** the original avatar 25 gram; **(B)** for the avatar with 20 times more mass than the original avatar without the passivity controller, and **(C)** with the passivity controller.

being interacted with by 20. Figure 6.5.4(B) shows how the system oscillates following the same user movement. Although the oscillation is bounded by physical damping from the user's hand, it can become unstable if the user releases the handle. The system kinetic energy peaks at $4.5Ncm$ then gradually decreases.

Passivity Control: M=20X In Figure 6.5.4(C), with the PCr active with a large mass, the system overshoots by 44% but converges within 200ms to the desired displacement. System energy peaks at 2.7Ncm and decreases more quickly than in the VC case for the same mass (B).

Example 2: User Interacts with a Virtual Mass-Spring System

The previous example showed how PCr can handle a large change in the system's virtual *mass*; how does it do with comparable changes in rendered stiffness as well as the same 20X mass range? We implement the system as illustrated in Figure 6.6, where a user draws a mass attached to a spring. Here, Figure 6.6(B) shows a graphical representation of the recognized mode. Our system recognizes a zigzag stroke as a spring and rectangle as a mass, and their connection on the sketch as a kinematic connection between them. The experience is similar to pulling on a real spring: force increases as one pulls further. Figure 6.6(B) shows the interaction result: as the user pulls down on the spring (change in X_{user}) by around 3cm and then "drops" the force – *i.e.*, stops resisting the haptic display's applied force – the system applies up to 1.27N of force to restore X_{user} to its starting position. The system exhibits a damped oscillation, with two sources: (a) the user's hand and (b) frictions in the haptic display. Here, this is desired behavior faithful to the virtual system dynamics in interaction with the user's hand damping, not a controller instability.

The graphical representation could optionally be displayed to the user to confirm recognition, and animated as they interact. Drawing and animating could implemented on a co-located tablet screen under the haptic display. In future we plan to investigate impacts of employing the user's original strokes versus changing them with a cleaner graphical representation, and of animating the diagrams.

The second row in Figure 6.6 shows the user placing two springs in parallel. The learning concept is that springs in parallel sum to a greater combined stiffness than in series, and the operator should feel a tighter (stiffer) connection. In comparison the previous example, the user should perceive a difference in force for the same displacement: the system supplies up to 2.3N force to the user's hand for a similar displacement to the single-spring case. As these results show, this system remains stable under passivity control for a doubling of total stiffness in combination with an already-large virtual mass. This mass-spring example can trivially be extended to include a damper (dissipative element). This is energetically less demanding – a virtual damper does not store energy. In general increasing virtual damping (assuming adequate sampling) reduces susceptibility to large impedance variation [82].

6.6 Discussion

We examine this work's contributions, and discuss how the PAL approach can be validated and extended.

6.6.1 The PAL Framework, Guidance and Exposed Needs

We drew on general theories of experiential learning to propose a framework that to help haptic and educational experts work together to leverage physical intuition in an effective learning process. This endeavor needs support: learning technology is notoriously hard to evaluate for efficacy, and get feedback on what is helpful. Despite evidence for the role of physical intuition and embodiment in effective learning, we know far less about how to saliently recreate it in digital contexts. Thus, rather than trying to show directly that haptic feedback helps learning, we built on a proven approach in first (a) accepting that *designing* and *exploring* are powerful supports to learning, then (b) seeing how haptic environments can make these activities more powerful than without them.

Metrics

While we have not yet evaluated our technical demonstrations with students, we will in future choose metrics (as per PAL-inspired goals) to highlight how the activities can be more fluid, engaging, focused, intuitive and insightful than without haptics.

Guidelines for PAL-Inspired Systems

In applying PAL principles we exposed some key requirements. We made progress in translating these to technical challenges, some of which can be addressed with current state-of-art techniques, and others where we need to further innovate. Here we summarize these, noting that while we have identified one pathway to implement them here (6.6.2), we hope that others will find more.

1. Let learners design their own worlds: PAL (and experiential learning theory generally) indicates that we should lower friction in letting learners (or in some case their teachers) build their environments. This is an old idea – Scratch and its ilk have born rich fruit – but we need this for environments amenable to haptic display for the purpose of accessing physical tuition.

2. Let learners explore, iterate and compare those worlds with physical feedback: Exploration should be informative, flexible and fun. Haptic feedback needs to be clear enough to support insights; it must be possible to jump around easily within an environment and try different things; and the whole process should flow, show insights that might not be otherwise available, surprise and delight. This entails a certain quality of haptic display, and curation of environments (*e.g.*, mechanical systems, electrical, hydraulic, chemistry) that while offering broad scope, also guide the learner on a rewarding path.

3. Moving between designing and exploring and back should be fluid: When experiential learning is working as it should, learners will generate more questions as they explore, and want to go back, re-design, compare and ask again. If they have to change modalities or undergo a laborious process to alter the environment or compare different examples, this cycle will be inhibited. We wonder if it is worth trying to stay (graphically) on paper while the digital world plays out through the haptic device, for immersion, focus and the intuitiveness of physical drawing; instead of fussing with a GUI.

4. Support a Broad Space for Experimentation: Instability is a continual risk for haptic force feedback systems, and could quickly turn anyone off as well as obscuring recognizable physical insights. Tightly restricting the explorable parameter space is an unacceptable solution, since it likewise limits the kinds of experiments to be conducted. Passivity control is one approach to a broader range than the methods currently available to novice hapticians via libraries.

6.6.2 Technical Proof-of-Concept

In the scope of this chapter, we have demonstrated at least one full technical pathway for a system that allows a user to design a haptically enabled system by sketching it on paper while adhering to some basic conventions, then interact with that system haptically – and stably – without changing mode or context across a parameter range of which is larger than typically supported in haptic environments. Its and-stroke recognition supports low-friction *designing*, so users can informally sketch ideas, even alter them. For *exploring,* we identified the inadequacy of the conventional rendering method of virtual coupling given the range of system parameters we need to support, and showed how a more specialized controller (based on passivity theory) could take it to this needed level. We encourage curators of haptic libraries to include passivity control support.

6.6.3 Generalizing to Other Physics Environments: a bond graph-Inspired Approach

Our examples demonstrate the ability of a passivity controller to bound a system's energy and prevent instability across a broad range of simulated system parameters. We did this based on a basic mechanical dynamic system, a mass oscillating with different spring combinations. This step can be translated with relative ease to other systems of interest in science learning.

Bond graph theory [109, 177] relates physical domains (*e.g.*, mechanics, electronics, hydraulics) based on energetic concepts of *efforts* and *flows*. This commonality is a means to connect domains, and also allows the translation of ideas between them. For our purposes, a physical model developed to represent a mechanical system can be translated with relative ease to an electrical domain.

Bond graphs hold threefold value here. First, technically we can exploit its analogies and representation to translate models and their support to other physical domains. Comparable properties will be relevant. In bond graphs, springs (mechanical) and capacitors (electrical) are analogs, both idealized to store energy in the same way, as are mass and inductance, dampers and resistors. Table 6.2, drawn from [24], includes a full list of bond domain analogies.

Second, these analogs provide a language and convention by which to render physical properties haptically: *e.g.*, effort, flow, resistance and inertance can be developed once and re-used in their relation to one another. It simplifies implementation in new domains.

Table 6.2: Analogy between some conventional physical domains, reproduced from Borutzky's [24].

Domain	**Flow**	**Effort**	**Compliance**	**Resistance**	**Inertance**
Electric	Current	Voltage	Capacitor	Resistor	Inductor
Kinetic translation	Velocity	Force	Spring	Damper	Mass
Kinetic rotational	Angular Velocity	Torque	Torsional Spring	Damper	Inertia
Hydraulic	Flow rate	Pressure	Chamber	Valve	Fluid inertia

Thirdly and most interesting pedagogically, these analogs are a powerful way to grasp and generalize fundamental relationships in physical systems. The haptic representation will reinforce this bond-centered generalization, helping learners to transfer their growing knowledge across domains: once they have mastered how the relations between current, voltage, compliance and resistance work in the electrical domain, they should be able to quickly apply them to kinetic or hydraulic systems. It is often the case that a learner feels more comfortable in one domain; they can use this 'home' grounding to support their understanding elsewhere.

6.7 Conclusion

A long-awaited promise of ubiquitous computing [223] is natural access to computational power where and when we need it. Yet, for the most part we remain tied to a small screen and a keyboard or tablet, with constrained space to work, keystroke input, a single viewport with many distractions, and interaction generally on the terms of the device.

In this chapter we proposed an approach to support multimodal learning with potential benefits to embodied learning and thinking. It includes a framework drawn from validated theories of experiential learning translated to the physical domain to guide system designers in creating educational systems focused on *designing* and *exploring*; underscoring of the importance of fluid, same-modality movement between these learning phases; demonstrations of the technical feasibility of implementing both idea capture and physical rendering in a pen-and-paper environment; and guidelines and assessment of how to move such a vision forward. We demonstrated these ideas on a fixed small-workspace device, but untethered, infinite workspace grounded force feedback has been prototyped and could be commercially viable given demand.

The present work points to a path away from tethered, disembodied interaction, examining ways to harness the natural fluidity and ease of pen-and-paper interactions and connect them to powerful digital simulation for the purpose of simulation, gaining physical, embodied insight, problem solving and thinking with our sense of touch as well as our heads and eyes. A graphical viewport is not always needed when we have our imagination, a sketchpad and hands to feel.

Chapter 7

Conclusions

We envision that a versatile digital manipulative should be able to support design and explore learning activities in response to continuous interactions with the learner. Below we discuss our contributions towards this vision.

7.1 Thesis Objectives and Contributions

In contrast to other excellent research in educational haptics, the work in this dissertation is shaped by an aspiration to develop and validate a versatile tool that empowers learners to perform operations on the object of knowledge. Based on Piaget's cognitive development theory, to know an object means to act on it. Operation as an essence of knowledge requires the learner to modify and transform an object, and also understand the process of transformation which leads to the knowledge of how the object is constructed [180]. We further evaluate whether doing specific operations with our tool can improve the way that users learn and think.

When designing an educational manipulative, whether a tangible interface or a robot, we observed that several researchers focused only on exploration. They often study learning impact by placing learners inside a simulation. The hypothesis in these studies is usually that learning will improve in a multi-modal learning environment as compared to using a single modality (often visual). Of course, some learners could bridge the gap between their previous knowledge to the new knowledge and learn the new concept. But we believe we can increase the chances for successful assimilation/accommodation by empowering learners to construct this bridge based on their hypothesis.

Design is the key element of Constructivism. Through design, learners actively construct their own knowledge. As opposed to expressing the results to learners and ask them to accept them, design allows learners to build upon their previous knowledge, make assumptions, come up with a hypothesis, and then evaluate the results. We did not find an educational tool that could aid learners during the process of designing and exploring; therefore, we sought a new digital manipulative that can help in both construction and exploration of new knowledge by exploiting physicality and ubiquitous computing.

A long-awaited promise of ubiquitous computing [223] is natural access to computational power where and when we need it. Yet, for the most part, we remain tied to a small screen and a keyboard or tablet, with constrained space to work, keystroke input, a single viewport with many distractions, and interaction generally on the terms of the device.

In this dissertation, we presented the design evolution and two applications (assisted sketching design and haptic exploration) for a novel digital manipulative in an educational context. It documents technical feasibility for a basic implementation of a pen-and-paper interaction approach to interactive, self-driven, exploration-centered physical simulation for the sake of learning and gaining physical insight about ideas. Much work remains before we can claim that the concept is ready for roll-out to students and teachers. In this chapter, we will summarize our contributions as well as our research findings, then outline the future steps towards our vision of MagicPen for learning.

7.1.1 Objective I: Design, Interaction Space, and Applications of a Low-Cost and Large Workspace Haptic Display

How can we create a low-cost, large workspace force feedback device? What type of new interactions can we support with it? What are the potential educational applications of this platform?

We made a low-cost, robust, and highly portable haptic stylus that can support two types of interaction: (a) force feedback assisted drawing, (b) haptic rendering of virtual environments. In Chapter 3, we introduce a novel, low-cost grounded force feedback device with an unlimited 2D workspace.

Our first contribution (Contribution I) was significant given the inaccessibility of current haptic displays due to the cost and workspace. For example, Haply [67] – a commercialized pantograph – costs about $300, despite its small workspace, which is insufficient for most school-environment applications. We expect the retail cost for MagicPen to be around $150 – $200 [1].

As we proceeded to optimize our design, we had to choose priorities. We focused on large strokes and fast communication as we thought they were more necessary for the type of interactions needed for designing and exploring. Future studies can optimize the MagicPen for better accuracy, higher amplitude and more resolution of the force feedback.

There are also opportunities to optimize our system both in cost and ergonomy. For instance, it is feasible to employ only two motors in the ballpoint drive, which would potentially reduces the cost even further and improves the ergonomy of the design of the stylus-based haptic device. Therefore we believe that the introduction of MagicPen lowers the barriers of entry for haptics in the education settings. We identified three primary types of haptic feedback that MagicPen is able to support, namely, 2D spatial guidance, 2D virtual fixtures, and vibrations. We only explored the interactions that require 2D force feedback, *e.g.*, navigational guidance, and virtual walls. The interactions that fully or partially rely on vibrotactile feedback and its combination with force feedback remain for future work.

Creating a functional platform was critical to the success of this dissertation and at the same time, had the highest risk among the rest of the objectives. We did several iterations to ascertain the saliency and consistency of the generated force

[1] Electronic components currently dominate cost, and can be reduced 40% through integration and more optimized choices, leading a reduction in parts cost from the current prototypes $100 to $60. For 500 samples we have:

- Device parts: $60 USD
- Assembly labour: $15 USD per device
- Marketing + Website + etc: $5000 USD
- Benefits:$10 – $30 USD per device
- Injection Moulding: $10000 – $15000 USD

Total cost= $115 – $155 USD per device= $148 CAD – $200 CAD feedback. The risk of this objective was high due to the fact that the results of the other objectives highly depended on the performance of our system. For instance, in case of failure, we needed to be able to differentiate between causes of no effect of haptics or low haptic quality.

MagicPen opens up new opportunities for a variety of applications including education, gaming, and assistive technologies. In this dissertation, we focused only on the educational class of applications. Instead of focusing on a particular educational problem or need, we studied how our device can impact the ubiquitous pen-and-paper interactions. Therefore, any positive result not only justifies the usefulness in a specific task but also pertains to a large impact size by taking into account the number of times it is being used throughout the learning process. Further investigations can check the usefulness and efficacy of MagicPen on specific learning disorders, *e.g.*, dyslexia.

7.1.2 Objective II: Phasking and Computer-Aided Design

What are the core interaction concepts for physically assisted sketching (phasking), and how we can support them with our force feedback pen?

While a major objective of the development of this technology was to support learning, usually by children or youth, we did not design MagicPen to be limited to a certain age range or be used just by children. If children see that a tool is being used by adults, they will not consider it as a toy and be more motivated to learn how to use it. Our Phasking application is an example that covers the full spectrum of users ranging from novices to skilled drawers. MagicPen not only helps designers who are less proficient at drawing to enhance their rapid sketches on paper but also aids experts to exploit their drawing skills in their CAD designs.

In our second contribution (Contribution II), we introduced *Bring, Bound, Control sharing, Tool selection, and Constructing constraint environment* as core interaction concepts in our Phasking framework. The details of how to support users highly depend on the application and require further in-depth investigation. We review the benefits of each of these proposed core interaction concepts in different sketching scenarios.

Constraints (Bounding and Bringing) – Using a traditional (familiar) set of tools (*e.g.*, ruler, compass and protractor) can aid both novice and experienced designers in creating more consistent and comprehensible sketches. However, tools (or the lack of the right one) can also slow them down. Our device offers both *bring* and *bound* constraints covering the whole variety of drawing assistants which the drawing assist tool-sets can bring to the table. Our results suggest how our system improves the accuracy of users' drawings in different tasks even when the users were putting the objects in perspective.

Tool Selection – MagicPen enhances communication by empowering a user's rapid sketching and drawing ability. We show how a user can select a specific functionality from a paper tool palette and accordingly the MagicPen helps them to draw the selected option. For instance, a user can select basic geometries and draw them with the help of MagicPen and then use them as the foundation for more complex drawings. In the current design, a user needs to explicitly select a tool from a paper tool palette. Another possibility is to implement more implicit physical assistants by predicting the user's intentions and help the user on the go.

Control Sharing – Our shared control drawing concept tries to preserve the creativity in assisted drawing by bringing the authority control to the hand of a user; therefore, the user can decide on the amount of assistant they receive. We hope this approach leads to more free collaboration between a user and the computer. Further studies are needed to elaborate on the use of control sharing in the creation of new art and assess the novelty and creativity aspects of it.

Constructing constraint environment – A key to expressive drawing is to construct and track the constraints. A user can define and set the constraint within the environment they operate. MagicPen supports the construction and required force feedback assistance in both digital and manual drawing mediums. To date, manual drawing and CAD drawing are performed in separate worlds, with few conversion options to move from one to the other.

We identified a gap between digital and manual drawing design spaces and we sought a new solution to close the gap. Our MagicPen uses built-in CAD software and is the first step towards a medium to interactively exchange pen-and-paper drawing information between a user and a computer – a process that is known as digital twinning.

Our work did not evaluate digital twinning and neither did we study how bringing ubiquitous computing directly to hand could enhance a designer's effectiveness and productivity. Future studies can investigate how these core interactions along with ubiquitous computing can impact the analysis, modification, or optimization of a design.

7.1.3 Objective III: Intuitive Learning of STEM with Haptics

How can we improve the versatility of the device through force feedback- the capacity to express information to users through the addition of haptics?

We strove for a versatile device that would make it suitable for several learning scenarios. In the third objective, we tried to understand the importance of physicality in STEM and uncover the useful strategies of employing haptics in learning activities. We tried to create more conclusive results through the lens of a collaborative learning framework. Specifically, we searched for how individuals try to construct, negotiate and share meanings using force feedback during the grounding process.

We learned that haptic feedback works as an unobtrusive channel to provide essential physical information to the user, to reinforce the visual cues by making a multi-modal experience, and to increase awareness of partners' presence and actions. It can also center one or a group of users' attention around a certain manipulation location, and lower cognitive effort through unobtrusive GUI-provided constraints, particularly as the GUI changes due to partner activity.
In our third contribution (Contribution III), we sought strategies of when and how force feedback should be offered to students to enhance their science learning reveals a large frequency of concurrence between exploration and evaluation of a hypothesis during the monitoring and reflection grounding acts. We also observed a similar trend of co-occurrence which suggests potential correlation between physical collaboration and the construction of a new environment to introduce a new topic.

Our closer look at the critical haptic events for three different environments paired with different haptic devices demonstrates how haptics can impact different dimensions of collaboration. Our results show a significant correlation between using haptics and reaching consensus across three environments. It also suggests a strong relationship between the use of haptics and sustaining mutual understanding as well as dialogue management. We observed different collaboration dynamics for the same dyad in three different environments suggesting that types of force feedback, the learning scenarios, and background knowledge, could impact the effectiveness of using haptics for learning STEM.

We shared the lessons we learned to help haptic designers and educational specialists for future research is touch/haptic sensory feedback for education to improve the chances of delivering a more successful haptic learning experience.

7.1.4 Objective IV: Physically assisted learning (PAL)

What are the key haptic interactions that can support learners throughout different stages of experiential learning cycle?

We connected the activities in design (Objective II) and exploration (Objective III) to achieve a smooth transition between these two stages in learning. We considered Kolb's [122] four stages of experiential learning and identify the key haptic interaction in each stage by reflecting on our findings from Chapter 4 and Chapter 5.

We introduced the Physically Assisted Learning (PAL) framework by focusing on the useful haptic interactions in two out of four stages of experiential learning i.e., Active Experimentation & Concrete Experience. More specifically, we focused on creating a smooth transition between these two stages through haptic augmentation by drawing the haptic experience and then feeling it.

We proposed an approach to instantiate this framework for two learning scenarios: a) handwriting assistant and b) mass-spring experiment. In both scenarios, the user creates the haptic experiences just by drawing them and then starts exploring. Accordingly, we unveil hidden technical difficulties as the haptic experiences levels up to more advanced renderings. For example, we addressed different challenges in stroke recognition or maintaining control stability in a highly dynamic situation. Finally, we offered a path forward to extend our approach to other domains such as physics and math.

As a part of Contribution III, we presented the theoretically-grounded PAL framework and demonstrate the technical feasibility of it. We found several pieces of evidence in related work to support our proposed framework; however, we still lack a user evaluation to confirm the manner and degree to which it actually supports learning. Validating this framework will require a series of focused studies that empirically evaluate the added value of physicality in *design* and *explore* learning phases as well as fluidity in the transition between them. These studies also need to consider factors such as engagement, ownership of knowledge, self-efficacy, self-confidence, and self-paced learning.

We see an immense value in studying *design* and *explore* together. It is crucial to know the impact of haptics in each part individually, but at the same time, these two components complement and reinforce each other. Conducting *design* activities will help learners to set the assumptions and define important parameters that eventually improve the learners' awareness of the explored environment. On the other hand, the experiences in *explore* can enhance the following iterations in design. Therefore a summative assessment is needed to articulate the larger question of how PAL can support operation on the object of knowledge.

7.2 Limitations

We mentioned some specific limitations above in the context of each objective. Here we discuss the higher level limitations that we found in the path towards evaluation and wide-spread use of MagicPen in education.

7.2.1 Limitations in Assessing Quality of Drawing

We used accuracy as an objective measure to assess the benefits of phasking on the quality of drawing. Accuracy matters when the user has a clear image of the drawing from the beginning. This is often not the case as the user's sketch extensively evolves over the course of drawing. The lack of clear expected outcomes inhibits careful evaluation of any drawing assist tools, especially when we need to maintain authorship and originality of the artwork.

In order to support this application, implicit assistance offers another class of phasking's supports that we did not study in this dissertation. As the drawing progresses, the computer agent tries to predict the user's intended drawing and accordingly apply proper force feedback assistance to the user's hand. The implicit assistance together with control sharing create a unique Human-AI collaboration that potentially preserve creativity while enhancing the quality of drawing.

We identified two objective measures to assess the quality of drawing for implicit phasking support with the goal of *communication* [133][56]. We can improve expressivity by assisting users to create:

- (a) more realistically proportioned drawings,
- (b) drawings with high-level perceptual information.

Even with these studies, we will still be far from making any general claim about the impact of phasking on a more artistic drawing or improving the drawing skills of the users.

7.2.2 Evaluations are Preliminary

Our qualitative approach towards understanding the importance of haptics in learning had a limited number of participants. It is often challenging to find the right sample size to represent a given learner population. Several factors and covariances contribute when we try to assess the impact of haptics and physicality in learning (*e.g.*, cultural background, socioeconomic status, and mental ability). One way to address this problem is to use a large sample size from a variety of schools. Large sample size can increase the statistical power and uncover individual differences. This becomes important when if the educational tool can only benefit some percentages of learners. We can still call this tool effective as long as we identify that specific group of learners who benefit from this tool the most.

We did not investigate learner's prior tactile experience and how it can impact the efficacy of MagicPen comprehensively. This was partially due to the fact that our participants did not have previous experiences with haptic displays (there were cases where learners were familiar with physical manipulatives). A longitudinal study could explore the long-term educational impact of haptics, as well as shed lights on how learners gradually develop skills that enhance using the sense of touch for learning.

7.2.3 Small Library of Learning Activities

We used some specific learning activities that could highlight the strengths of our PAL framework. However, similarly to other educational robots, we need to have a variety of concrete lesson plans and a set of objectives to promote learning activities with MagicPen. This requires a substantial investment and more resources beyond what exists in our lab. Focusing on teachers as the main orchestrators in classrooms, we tried to simplify this process and lower the technical barriers for inviting teachers to design these learning activities, set learning objectives, and eventually guide learners towards them. However, more research is required to provide high incentives to teachers to get involved and adopt this technology in their classrooms.

In addition, we lack a community platform that teachers and learners can share their results and findings with their peers. This platform is where teachers can present their learning activities and share the lesson learned. Learners, on the other hand, can discuss their problems, learn from each other, and together tackle new problems. There is no need to pose any age or location constraint on this platform. We can create a learning society where learning occurs across learners' lifespans.

7.3 Future Work: The Path Forward

We lay out some foreseeable next steps.

Basic Access and Versatility

First and foremost, building haptically augmented worlds and even accessing them interactively requires considerable expertise and infrastructure. Haptic technology is anything but accessible, and this barrier will need to be breached. As for any educational technology, the principal barriers will be cost, robustness, versatility, and usability or expertise.

The vision which will eventually spur the needed technical refinement is *versatility*: myriad ways to use physical interaction in a form factor that one can carry around, perhaps first like an engineering calculator then becoming more ubiquitously useful as a haptically augmented smartphone stylus. Versatility is needed first in use case development, and following that in device form factor. In order to enhance MagicPen's versatility, we can incorporate other modalities to communicate richer information to the user. It can embed an E-ink LCD to graphically present information based on its location, or use a microphone to receive the user's command or understand their utterance and behave accordingly.

Usability and Expertise: A theme in this dissertation has been lowering friction and barriers to entry for both teachers and learners. This also needs to become more true for system designers, allowing them to participate in system development from their home discipline and even without engineering expertise – *e.g.*, education experts. Input methods, library construction, support groups, and other aspects of development ecosystems will move us in this direction.

Eventually, Logistical Deployment with Kids: Classrooms are challenging environments; just batteries, power cords and updating host computers present almost insurmountable obstacles. The first point of contact may be science centers and tutoring centers, and potentially on to personal devices (like student calculators) rather than school-supplied technology. A key to classroom adoption is teachers' awareness about assistive learning technology in their classroom by monitoring the status of MagicPens in a dashboard. This enables the teachers to better orchestrate their classroom and for immediate troubleshooting purposes.

Enhanced Usability, Fluidity and Function

Throughout this dissertation, we have described many possible variations and augmentations to our basic implementation, all of which can be explored to discover optimality from logistic and pedagogical standpoints, and inform the direction of further technology development. To name a few (and going beyond innovation in the haptic technology itself):

- CAD-type sketching support at the *design* stage
- More advanced sketch-recognition functions, *e.g.*, setting and modifying simulation parameter values by [re]writing them on paper
- Generating more extensive simulation environments, in multiple domains (*e.g.*, bond graph extensions)
- Utilizing more sophisticated haptic rendering algorithms as we encounter limits
- Finding good haptic representations for abstract fields such as maths
- Libraries to support educators setting up "design sandboxes".

7.4 Final Remarks

In this dissertation, we tried to make a vehicle to deliver Piaget's theory of Constructivism and used this lens to both innovate and study alternative forms of haptic technology. We learned that note-taking on paper is not just for information retention but are one of the main places where the construction of knowledge occurs. Paper interactions have already become smart through years of innovation in touch-enabled LCDs technology on tablets and laptops. So instead of paper, we focused on the pens. My supervisor's and my efforts can be summarized as adding force feedback to pens and evaluate the potential educational benefits. We tried to keep the device low-cost to be accessible by every motivated learner. The outcome turns into a fully standalone device that can be a powerful tool with minimal requirements including energy and a piece of paper. Besides these requirements, the *pick-and-play* feature of MagicPen potentially can lower the logistic barriers for entering the classrooms.

It is most likely that the virtual world will become a big part of our lives, and education will not be an exception. The field of haptics offers realism, immersion, and expressivity [118] to extend the physical experiences into the virtual world. This gives MagicPen a unique opportunity to push us toward mixed reality learning, where the learning does not stop in one world. Removing XR goggles puts an end to the users' interaction with the virtual world. Unlike XR goggles, MagicPen will be with the user throughout the whole experience. We call it a "haptic ecotone" between the physical and the virtual world – a learning companion that carries the experiences forward.

Our final note is essentially a reiteration of this dissertation's objectives. Beyond the creation of MagicPen itself, we learned that we should let the learners be in charge of constructing and exploring their learning environments. As a result, learners are able to hypothesize and then critique their own thinking, as opposed to traditional methods which expose learners to a pre-designed phenomenon in an explanatory manner. We should not confine learners to the designer's imagination or mental model of the world. For example, learners cannot try out the effect of negative gravity unless this condition has been foreseen and implemented by the designer. Arguably, we believe that this contradiction arises from the Post-positivism approach that we are taking to simulate the world in a virtual environment versus the

Constructivism requirements we need to create a hypothetical world for education. For instance, one of the key aspects of a Constructivism's learning environment is the freedom to make mistakes and learn from them. However, this freedom comes at a cost, when learners do not know what exactly to do and therefore need guidance. The challenge is to provide just enough guidance to lead the learner in the right direction without limiting them to the designer/instructor's imagination.

Bibliography

[1] 3Dsystems. *Touch haptic device*, accessed February 15, 2019. URL https://www.3dsystems.com/haptics-devices/touch. → pages 4, 23

[2] D. Abbink and M. Mulder. Neuromuscular analysis as a guideline in designing shared control. In M. H. Zadeh, editor, *Advances in Haptics*, pages 23–35. IntechOpen, 2010. → pages 19

[3] D. A. Abbink, M. Mulder, and E. R. Boer. Haptic shared control: smoothly shifting control authority. *Journal of Cognition, Technology and Work*, pages 19–28, 2012. → pages 19, 52, 55

[4] H. Abelson and A. A. DiSessa. *Turtle geometry: The computer as a medium for exploring mathematics*, volume 8. MIT press, 1986. → pages 2

[5] E. K. Ackermann. 26 experiences of artifacts. *Designing Constructionist Futures: The Art, Theory, and Practice of Learning Designs*, page 275, 2020. → pages 124, 127, 129

[6] B. D. Adelstein and M. J. Rosen. Design and implementation of a force reflecting manipulandum for manual control research. *Advances in Robotics*, pages 1–12, 1992. → pages 16

[7] S. Adhikari. Stylus based haptic peripheral for touch screen and tablet devices, 2011. URL https://patents.google.com/patent/US8681130. US Patent US8681130. → pages 22

[8] M. M. Amin, H. B. Zaman, and A. Ahmad. Visual haptic approach complements learning process of jawi handwriting skills. In *2013 5th International Conf on Information and Communication Technology for the Muslim World (ICT4M)*, pages 1–6, 2013. → pages 135

[9] A. N. Antle and A. F. Wise. Getting Down to Details: Using Theories of Cognition and Learning to Inform Tangible User Interface Design. *Interacting with Computers*, pages 1–20, 2013. → pages 123, 126

[10] T. Asselborn, A. Guneysu, K. Mrini, E. Yadollahi, A. Ozgur, W. Johal, and P. Dillenbourg. Bringing letters to life: Handwriting with haptic-enabled tangible robots. In *Proceedings of the 17th ACM Conf on Interaction Design and Children (IDC)*, pages 219–230, 2018. → pages 120, 135

[11] S. Atmatzidou, S. Demetriadis, and P. Nika. How does the degree of guidance support students' metacognitive and problem solving skills in educational robotics? *Journal of Science Education and Technology*, pages 70–85, 2018. → pages 124

[12] C. A. Avizzano, M. Satler, G. Cappiello, A. Scoglio, E. Ruffaldi, and M. Bergamasco. Motore: A mobile haptic interface for neuro-rehabilitation. In *Robot and Human Interactive Communication (RO-MAN)*, pages 383–388, 2011. → pages 27

[13] M. Baker, T. Hansen, R. Joiner, and D. Traum. The role of grounding in collaborative learning tasks. *Journal of Collaborative Learning: Cognitive and Computational Approaches*, 0:63–75, 1999. → pages xx, 73, 75, 79, 82

[14] B. Barron. Problem solving in video-based microworlds: Collaborative and individual outcomes of high-achieving sixth-grade students. *Journal of Educational Psychology*, pages 391–405, 2000. → pages 78

[15] A. L. Barrow and W. S. Harwin. High bandwidth, large workspace haptic interaction: Flying phantoms. In *2008 Symp on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 295–302, 2008. → pages 17

[16] O. Bau, I. Poupyrev, A. Israr, and C. Harrison. Teslatouch: Electrovibration for touch surfaces. *Proceedings of the 23nd annual ACM on User interface software and technology*, page 283–292, 2010. → pages 25

[17] M. Baumann, K. E. MacLean, T. W. Hazelton, and A. McKay. Emulating human attention-getting practices with wearable haptics. In *IEEE Haptic Symp (HAPTICS)*, pages 149–156, 2010. → pages 72

[18] B. Bayart, A. Pocheville, and A. Kheddar. An adaptive haptic guidance software module for i-touch: example through a handwriting teaching simulation and a 3d maze. In *IEEE International Workshop on Haptic Audio Visual Environments and their Applications*, page 6, 2005. → pages 48

[19] N. Beckers, E. H. van Asseldonk, and H. van der Kooij. Haptic human–human interaction does not improve individual visuomotor adaptation. *Scientific reports*, pages 1–11, 2020. → pages 128

[20] P. Berkelman, B. Tix, and H. Abdul-Ghani. Electromagnetic position sensing and force feedback for a magnetic stylus with an interactive display. *IEEE Magnetics Letters*, pages 1–5, 2019. → pages 25

[21] A. K. Bhunia, A. Das, U. R. Muhammad, Y. Yang, T. M. Hospedales, T. Xiang, Y. Gryaditskaya, and Y.-Z. Song. Pixelor: A competitive sketching ai agent. so you think you can sketch? *ACM Trans. Graph.*, pages 166–181, 2020. → pages 132

[22] S. Blanchard, V. Freiman, and N. Lirrete-Pitre. Strategies used by elementary school children solving robotics-based complex tasks: Innovative potential of technology. *Procedia-Social and Behavioral Sciences*, pages 2851–2857, 2010. → pages 124

[23] D. Bloom. Collaborative test taking: Benefits for learning and retention. *College Teaching*, pages 216–220, 2009. → pages 78

[24] W. Borutzky. *Bond graph modelling of engineering systems*. Springer Science Business Media, 2011. → pages xvi, 150

[25] BOSCH. Absolute orientation sensor. 2015. URL https://www.bosch-sensortec.com/bst/products/all_products/bno055. → pages 39

[26] V. Braun and V. Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, pages 77–101, 2006. → pages 99

[27] S. Brave and A. Dahley. intouch: A medium for haptic interpersonal communication. In *Conf. on Human Factors in Computing Systems (CHI)*, pages 363–364. ACM Press, 1997. → pages 72

[28] K. A. Bruffee. Collaborative learning and the conversation of mankind". *College English*, pages 635–652, 1984. → pages 78

[29] G. Burdea, J. Zhuang, E. Roskos, D. Silver, and N. Langrana. A portable dextrous master with force feedback. *Presence: Teleoperators and Virtual Environments*, pages 18–28, 1992. → pages 16

[30] D. Cesareni, S. Cacciamani, and N. Fujita. Role taking and knowledge building in a blended university course. *International Journal of Computer-Supported Collaborative Learning*, pages 9–39, 2016. → pages xv, 91

[31] A. Chan, K. E. MacLean, and J. McGrenere. Designing haptic icons to support collaborative turn-taking. *Int'l J Human Computer Studies*, 66: 333–355, 2008. → pages 73

[32] Y. Cho, A. Bianchi, N. Marquardt, and N. Bianchi-Berthouze. Realpen. *Proceedings of the 29th Annual Symp on User Interface Software and Technology*, pages 30–43, 2016. → pages 29

[33] S. Choi and K. J. Kuchenbecker. Vibrotactile display: Perception, technology, and applications. *Proceedings of the IEEE, Special Issue on Perception-Based Media Processing*, pages 2093–2104, 2013. → pages 16

[34] E. C. Chubb, J. E. Colgate, and M. A. Peshkin. Shiverpad: A glass haptic surface that produces shear force on a bare finger. *IEEE Trans on Haptics*, page 189–198, 2010. → pages 25

[35] H. H. Clark and E. F. Schaefer. Contributing to discourse. *Cognitive science*, pages 259–294, 1989. → pages 86

[36] M. M. Clifford. Thoughts on a theory of constructive failure. *Educational Psychologist*, pages 108–120, 1984. → pages 130

[37] T. R. Coles, D. Meglan, and N. W. John. The role of haptics in medical training simulators: A survey of the state of the art. *IEEE Trans on Haptics*, pages 51–66, 2011. → pages 23

[38] J. E. Colgate and J. M. Brown. Factors affecting the z-width of a haptic display. In *IEEE International Conf on Robotics Automation*, pages 3205–3210, 1994. → pages 143

[39] J. E. Colgate and G. G. Schenkel. Passivity of a class of sampled-data systems: Application to haptic interfaces. *Journal of Robotic Systems J. Robotic Syst.*, page 37–47, 1997. → pages 143

[40] J. E. Colgate, W. Wannasuphoprasit, and M. Peshkin. Cobots: Robots for collaboration with human operators. In *the 5th Ann. Symp. on Haptic Interfaces for Virtual Environment and Teleoperator Systems, ASME/IMECE*, volume DSC, pages 443–439, 1996. → pages 21, 43, 55

[41] F. Conti, F. Barbagli, R. Balaniuk, M. Halg, C. Lu, D. Morris, L. Sentis, J. Warren, O. Khatib, and J. K. Salisbury. The CHAI libraries. pages 496–500, 2003. → pages 18

[42] A. L. Cornu. Meaning, internalization, and externalization: Toward a fuller understanding of the process of reflection and its role in the construction of the self. *Adult Education Quarterly*, pages 279–297, 2009. → pages 124

[43] H. Culbertson and K. Kuchenbecker. Ungrounded haptic augmented reality system for displaying roughness and friction. *IEEE/ASME Trans Mechatronics*, pages 1839–1849, 2017. → pages 22, 26, 51

[44] R. Davis, M. Martina-Ortiz, O. Schneider, K. E. MacLean, A. Okamura, and P. Blikstein. The haptic bridge: Towards a theory of haptic-supported learning. In *Interaction Design and Children (IDC)*, pages 45–53, 2017. → pages 7, 28, 74, 77, 84, 129

[45] T. de Jong, M. C. Linn, and Z. C. Zacharia. Physical and virtual laboratories in science and engineering education. *Science*, pages 305–308, 2013. → pages 78

[46] J. Dewey. Logic: The theory of inquiry. 1938. → pages 123

[47] S. Dey, P. Riba, A. Dutta, J. Llados, and Y.-Z. Song. Doodle to search: Practical zero-shot sketch-based image retrieval. In *Proceedings of the IEEE/CVF Conf on Computer Vision and Pattern Recognition*, pages 2179–2188, 2019. → pages 132

[48] P. Dillenbourg, D. Traum, and D. Schneider. Grounding in multi-modal task-oriented collaboration. In *Proceedings of the European Conf on AI in Education*, pages 401–407, 1996. → pages xv, 73, 86, 87, 89, 90

[49] N. Diolaiti, G. Niemeyer, F. Barbagli, and J. K. Salisbury. Stability of haptic rendering: Discretization, quantization, time delay, and coulomb effects. *IEEE Trans on Robotics*, pages 256–268, 2006. → pages 132

[50] D. Dixon, M. Prasad, and T. Hammond. icandraw: Using sketch recognition and corrective feedback to assist a user in drawing human faces. In *Proceedings of the SIGCHI Conf on Human Factors in Computing Systems*, pages 897–906. ACM, 2010. → pages 51

[51] S. Do-Lenh, P. Jermann, A. Legge, G. Zufferey, and P. Dillenbourg. Tinkerlamp 2.0: Designing and evaluating orchestration technologies for the classroom. In *21st Century Learning for 21st Century Skills*, pages 65–78. Springer Berlin Heidelberg, 2012. → pages 77

[52] W. Doise, G. Mugny, A. S. James, N. Emler, and D. Mackie. *The social development of the intellect*. Elsevier, 2013. → pages 78

[53] P. R. . Electronics. Micro metal gearmotors. 2019. URL https://www.pololu.com/category/60/micro-metal-gearmotors. → pages 57

[54] F. L. Engel, P. Goossens, and R. Haakma. Improved efficiency through i- and e-feedback: a trackball with contextual force feedback. *International Journal of Human-Computer Studies*, pages 949 – 974, 1994. → pages 24

[55] M. A. Evans. Rapid prototyping and industrial design practice: can haptic feedback modelling provide the missing tactile link? *Rapid Prototyping Journal*, 2005. → pages 128

[56] J. E. Fan, D. L. Yamins, and N. B. Turk-Browne. Common object representations for visual production and recognition. *Cognitive science*, pages 2670–2698, 2018. → pages 160

[57] M. Fan, A. N. Antle, and E. S. Cramer. Exploring the design space of tangible systems supported for early reading acquisition in children with dyslexia. In *Proceedings of the TEI '16: Tenth International Conf on Tangible, Embedded, and Embodied Interaction - TEI*, pages 689–692, 2016. → pages 126

[58] N. Fellion, T. Pietrzak, and A. Girouard. Flexstylus: Leveraging bend input for pen interaction. In *Proceedings of the 30th Annual ACM Symp on User Interface Software and Technology*, page 375–385, 2017. → pages 29

[59] P. Fernando, R. L. Peiris, and S. Nanayakkara. I-draw: Towards a freehand drawing assistant. In *Proceedings of the 26th Australian Computer-Human Interaction Conf on Designing Futures: The Future of Design*, pages 208–211, 2014. → pages 23, 51, 52, 53

[60] J. C. Flanagan. The critical incident technique. *Psychological bulletin*, page 327, 1954. → pages 97

[61] Forcedimension. *Omega high precision force feedback interface*, accessed February 16, 2019. URL http://www.forcedimension.com/products. → pages 4, 23

[62] J. Forsslund, M. Yip, and E.-L. Sallnäs. Woodenhaptics: A starting kit for crafting force-reflecting spatial haptic devices. In *Conf on Tangible, Embedded, & Embodied Interaction (TEI)*, pages 133–140, 2015. → pages 28

[63] B. Forsyth and K. E. MacLean. Predictive haptic guidance: Intelligent user assistance for the control of dynamic tasks. *IEEE Trans on Visualization and Computer Graphics*, pages 103–113, 2006. → pages 55

[64] P. Frei, V. Su, B. Mikhak, and H. Ishii. Curlybot: Designing a new class of computational toys. In *Proceedings of the SIGCHI Conf on Human Factors in Computing Systems*, pages 129–136, 2000. → pages 2, 6, 124

[65] C. Gallacher and S. Ding. The haply development platform: A modular, open-source haptic ecosystem that enables accessible multi-platform simulation development. In *IEEE Haptics Symp (HAPTICS)*, page 156, 2018. → pages 133

[66] C. Gallacher and S. Ding. *hAPI library*, February 2017 (accessed September 15, 2017). URL https://github.com/HaplyHaptics/hAPI. → pages 18, 133, 137

[67] C. Gallacher, A. Mohtat, S. Ding, and J. Kövecses. Toward open-source portable haptic displays with visual-force-tactile feedback colocation. In *IEEE HAPTICS*, pages 65–71, 2016. → pages xix, 23, 26, 28, 44, 74, 76, 81, 86, 154

[68] T. Gargot, T. Asselborn, I. Zammouri, J. Brunelle, W. Johal, P. Dillenbourg, D. Archambault, M. Chetouani, D. Cohen, and S. M. Anzalone. "it is not the robot who learns, it is me." treating severe dysgraphia using child–robot interaction. *Frontiers in Psychiatry*, page 5, 2021. → pages 135

[69] S. Gerofsky. Seeing the graph vs. being the graph: Gesture, engagement and awareness in school mathematics. In *Integrating gestures*, pages 10–18. 2011. → pages 28

[70] D. C. Giancoli. *Physics: principles with applications*, volume 4. Pearson/Prentice Hall Upper Saddle River, NJ, 2005. → pages 139

[71] B. Gillespie and M. Cutkosky. Stable user-specific haptic rendering of the virtual wall. In *ASME International Mechanical Engineering, Symp on Haptic Interfaces*, pages 397–406, 1996. → pages 136

[72] R. B. Gillespie and A. M. Okamura. Haptic interaction for hands-on learning in system dynamics and controls. *Control Systems Magazine*, pages 12–23, 2008. → pages 20

[73] A.-L. Godhe, P. Lilja, and N. Selwyn. Making sense of making: critical issues in the integration of maker education into schools. *Technology, Pedagogy and Education*, pages 317–328, 2019. → pages 3

[74] E. B. Goldstein. *Encyclopedia of perception*. Sage, 2010. → pages 74

[75] N. Goodman. *Ways of worldmaking*. Hackett Publishing, 1978. → pages 127

[76] P. Griffiths and R. B. Gillespie. Shared control between human and machine: Haptic display of automation during manual control of vehicle heading. In *the 12th International on Haptic Interfaces for Virtual Environment and Teleoperator Systems (HAPTICS)*, pages 358–366, 2004. → pages 52

[77] P. G. Griffiths and R. B. Gillespie. Sharing control between humans and automation using haptic interface: Primary and secondary task performance benefits. *Human Factors*, pages 574–590, 2005. → pages 20

[78] T. Grossman, R. Balakrishnan, G. Kurtenbach, G. Fitzmaurice, A. Khan, and B. Buxton. Interaction techniques for 3d modeling on large displays. In *ACM Symp Interactive 3D Graphics*, pages 17–23, 2001. → pages 35

[79] A. Guneysu Ozgur, A. Özgür, T. Asselborn, W. Johal, E. Yadollahi, B. Bruno, M. Skweres, and P. Dillenbourg. Iterative design and evaluation of a tangible robot-assisted handwriting activity for special education. *Frontiers in Robotics and AI*, page 29, 2020. → pages 135

[80] A. Gürel. Cognitive comparison of using hand sketching and parametric tools in the conceptual design phase. 2019. → pages 30

[81] A. K. H. Habash. Touch pen with haptic feedback, 2015. URL https://patents.google.com/patent/US9116560B1/en?oq=+9%2c116%2c560. US Patent US9116560B1. → pages 22

[82] A. Haddadi, K. Razi, and K. Hashtrudi-Zaad. Operator dynamics consideration for less conservative coupled stability condition in bilateral teleoperation. *IEEE/ASME Trans on Mechatronics*, pages 2463–2475, 2015. → pages 132, 147

[83] G. Hallman, I. Paley, I. Han, and J. B. Black. Possibilities of haptic feedback simulation for physics learning. In *EdMedia+ Innovate Learning*, pages 3597–3602, 2009. → pages 74, 126

[84] F. G. Hamza-Lup and I. A. Stanescu. The haptic paradigm in education: Challenges and case studies. *The Internet and Higher Education*, pages 78–81, 2010. → pages 128

[85] I. Han and J. B. Black. Incorporating haptic feedback in simulation for learning physics. *Computers Education*, pages 2281 – 2290, 2011. → pages 78

[86] B. Hannaford and J.-H. Ryu. Time-domain passivity control of haptic interfaces. *IEEE Trans on Robotics and Automation*, pages 1–10, 2002. → pages xxii, 143, 144

[87] Hapticshouse. *Falcon Haptic Devices*, accessed February 16, 2019. URL https://hapticshouse.com/. → pages 4, 23

[88] S. Harada, S. Saponas, and J. A. Landay. Voicepen: Augmenting pen input with simultaneous non-linguistic vocalization. In *Proceedings of the 9th International ACM Conf on Multimodal Interfaces*, pages 1–15, 2007. → pages 28

[89] V. Hayward and O. R. Astley. Performance measures for haptic interfaces. In *Robotics Research*, pages 195–206, 1996. → pages 19

[90] V. Hayward and K. E. Maclean. Do it yourself haptics: Part i. *IEEE Robotics Automation Magazine*, pages 88–104, 2007. → pages 35

[91] F. Hemmert, A. Müller, R. Jagodzinski, G. Wintergerst, and G. Joost. Reflective haptics: Haptic augmentation of guis through frictional actuation of stylus-based interactions. In *ACM Symp on User Interface Software and Technology (UIST)*, pages 383–384, 2010. → pages 26

[92] J. Holst-Wolf, Y.-T. Tseng, and J. Konczak. The minnesota haptic function test. *Frontiers in Psychology*, page 818, 2019. → pages 73

[93] P. Honey and A. Mumford. *The learning styles helper's guide*. Peter Honey Publications Maidenhead, 2000. → pages xxi, 121

[94] J. Huegel and M. O'Malley. Progressive haptic and visual guidance for training in a virtual dynamic task. *IEEE Haptics Symposium*, pages 343–350, 2010. → pages 135

[95] G. Huisman. Social touch technology: A survey of haptic technology for social touch. *IEEE Trans on Haptics*, page 391–408, 2017. → pages 72

[96] M. Hung, D. Ledo, and L. Oehlberg. Watchpen: Using cross-device interaction concepts to augment pen-based interaction. pages 1–8, 2019. → pages 29

[97] A. Inc. Anoto digital pen. 2017. URL http://www.anoto.com/. → pages 51

[98] N. S. Inc. Neo smartpen m1. 2018. URL https://www.neosmartpen.com. → pages 39, 51, 133

[99] A. Ioannou and E. Makridou. Exploring the potentials of educational robotics in the development of computational thinking: A summary of current research and practical proposal for future work. *Education and Information Technologies*, pages 2531–2544, 2018. → pages 130

[100] N. Jafari, K. D. Adams, and M. Tavakoli. Haptics to improve task performance in people with disabilities: A review of previous studies and a guide to future research with children with disabilities. *Journal of rehabilitation and assistive technologies engineering*, page 147, 2016. → pages 128

[101] M. Jalal al Din Rumi. Townspeople, who have never seen an elephant, examine its appearance in the dark. the walters art museum, 1663 AD (1073 AH). URL https://art.thewalters.org/detail/83750. → pages xix, 73

[102] Y. Jansen, T. Karrer, and J. Borchers. Mudpad: Localized tactile feedback on touch surfaces. In *Adjunct Proceedings of the 23Nd Annual ACM Symp on User Interface Software and Technology*, UIST, pages 385–386, 2010. → pages 26

[103] C. Johnson. *Harold and the Purple Crayon*. Harper Brothers, 1955. → pages 4, 34

[104] L. A. Jones and M. Berris. The psychophysics of temperature perception and thermal-interface design. In *Proceedings of the 10th Symp on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, page 137, 2002. → pages 114

[105] M. G. Jones, A. Bokinsky, T. Andre, D. Kubasko, A. Negishi, R. Taylor, and R. Superfine. Nanomanipulator applications in education: the impact of haptic experiences on students' attitudes and concepts. pages 279–282, 2002. → pages 128

[106] D. L. Kain. Owning significance: The critical incident technique in research. *Foundations for research: Methods of inquiry in education and the social sciences*, pages 69–85, 2004. → pages 96

[107] P. Kammermeier, A. Kron, J. Hoogen, and G. Schmidt. Display of holistic haptic sensations by combined tactile and kinesthetic feedback. *Presence*, pages 1–15, 2004. → pages 16

[108] M. Kapur and K. Bielaczyc. Designing for productive failure. *Journal of the Learning Sciences*, pages 45–83, 2012. → pages 130

[109] D. C. Karnopp, D. L. Margolis, and R. C. Rosenberg. *System dynamics: a unified approach*. 1990. → pages 150

[110] O. B. Kaul, M. Pfeiffer, and M. Rohs. Follow the force: Steering the index finger towards targets using ems. In *Proceedings of the CHI Conf Extended Abstracts on Human Factors in Computing Systems*, CHI, pages 2526–2532, 2016. → pages 51

[111] A. Kay. Tesseract: An open-source optical character recognition engine. *Linux J.*, (159):2, 2007. → pages 141

[112] H. Khodr, S. Kianzad, W. Johal, A. Kothiyal, B. Bruno, and P. Dillenbourg. Allohaptic: Robot-mediated haptic collaboration for learning linear functions*. In *2020 29th IEEE International Conf on Robot and Human Interactive Communication (RO-MAN)*, pages 27–34, 2020. → pages 75, 77, 78, 84, 120, 128

[113] S. Kianzad and K. E. MacLean. Harold's purple crayon rendered in haptics: Large-stroke, handheld ballpoint force feedback. In *2018 IEEE Haptics (HAPTICS)*, pages 106–111, 2018. → pages xix, 57, 61, 64, 76, 81, 86, 120

[114] S. Kianzad and K. E. MacLean. Collaborating through magic pens: Grounded forces in large, overlappable workspaces. In *Haptic Interaction*, pages 233–237, 2019. → pages vii, 84, 121

[115] S. Kianzad, S. O. Karkouti, and H. D. Taghirad. Force control of intelligent laparoscopic forceps. *Journal of Medical Imaging and Health Informatics*, pages 284–289, 2011. → pages 114

[116] S. Kianzad, Y. Huang, R. Xiao, and K. E. MacLean. Phasking on paper: Accessing a continuum of physically assisted sketching. In *ACM Conf on Human Factors in Computing Systems (CHI)*, pages 1–11, 2020. → pages 81, 120, 128, 135

[117] S. Kianzad, G. Chen, and K. E. MacLean. PAL: A framework for physically assisted learning through design and exploration with a haptic robot buddy. *Frontiers in Robotics and AI*, 8:298–250, 2021. → pages vii

[118] E. Kim and O. Schneider. Defining haptic experience: Foundations for understanding, communicating, and evaluating hx. *Proceedings of the CHI Conf on Human Factors in Computing Systems*, page 1–13, 2020. → pages 163

[119] N. W. Kim, N. Henry Riche, B. Bach, G. A. Xu, M. Brehmer, K. Hinckley, M. Pahud, H. Xia, M. McGuffin, and H. Pfister. Datatoon: Drawing dynamic network comics with pen + touch interaction. In *CHI*, pages 1–12, 2019. → pages 31

[120] R. L. Klatzky, D. Pawluk, and A. Peer. Haptic perception of material properties and implications for applications. *Proceedings of the IEEE*, pages 2081–2092, 2013. → pages 16

[121] D. L. Knee. Dynamic resistance control of a stylus, May 3 2007. URL https://patents.google.com/patent/US7508382B2/en?oq=US+Pat.+No+7%2c508%2c382+B2. US Patent US7265750B2. → pages 22

[122] D. Kolb. *Experiential learning: experience as the source of learning and development*. Prentice Hall, 1984. → pages xxi, 118, 120, 121, 125, 158

[123] C. Kontra, D. J. Lyons, S. M. Fischer, and S. L. Beilock. Physical experience enhances science learning. *Psychological science*, pages 737–749, 2015. → pages 78

[124] M. Konyo. Tako-pen: A pen-type pseudo-haptic interface using multipoint suction pressures. In *SIGGRAPH Asia 2015 Haptic Media And Contents Design*, pages 2–4, 2015. → pages 29

[125] G. Korres and M. Eid. Katib: Haptic-visual guidance for handwriting. In *International Conf on Human Haptic Sensing and Touch Enabled Computer Applications*, pages 279–287. Springer, 2020. → pages 135

[126] U. Kuckartz and S. Rädiker. *Analyzing qualitative data with MAXQDA*. Springer, 2019. → pages 91

[127] K.-U. Kyung, J.-Y. Lee, and J. Park. Design and applications of a pen-like haptic interface with texture and vibrotactile display. In *2007 Frontiers in the Convergence of Bioscience and Information Technologies*, pages 543–548, 2007. → pages 128

[128] J. T. F. Laurent Denoue. Force-feedback stylus and applications to freeform ink, May 3 2005. URL https://patents.google.com/patent/US7508382B2/

en?oq=US+Pat.+No+7%2c508%2c382+B2. US Patent US7508382B2. → pages 22

[129] T. B. Lauwers, G. A. Kantor, and R. L. Hollis. A dynamically stable single-wheeled mobile robot with inverse mouse-ball drive. In *IEEE Int'l Conf on Robotics and Automation (ICRA)*, pages 2884–2889, 2006. → pages 27, 38, 41

[130] J. J. LaViola and R. C. Zeleznik. Mathpad 2: A system for the creation and exploration of mathematical sketches. In *ACM SIGGRAPH*, page 432–440, 2004. → pages 30

[131] S. J. Lederman and R. L. Klatzky. Hand movements: A window into haptic object recognition. *Cognitive Psychology*, pages 342–368, 1987. → pages 73, 74

[132] J. C. Lee, P. H. Dietz, D. Leigh, W. S. Yerazunis, and S. E. Hudson. Haptic pen: A tactile feedback stylus for touch screens. In *Proceedings of the 17th Annual ACM Symp on User Interface Software and Technology*, UIST, pages 291–294. ACM, 2004. → pages 23

[133] Y. J. Lee, C. L. Zitnick, and M. F. Cohen. Shadowdraw: Real-time user guidance for freehand drawing. *ACM Trans. Graph.*, page 1–27, 2011. → pages 51, 160

[134] V. Levesque, L. Oram, K. MacLean, A. Cockburn, N. D. Marchuk, D. Johnson, J. E. Colgate, and M. A. Peshkin. Enhancing physicality in touch interaction with programmable friction. In *Proceedings of the SIGCHI Conf on Human Factors in Computing Systems*, CHI, pages 2481–2490. ACM, 2011. → pages 25

[135] G. Levin. *Painterly Interfaces for Audiovisual Performance*. PhD thesis, Massachusetts Institute of Technology, Cambridge, 2000. → pages 49

[136] Y. Li, J. C. Huegel, V. Patoglu, and M. K. O'Malley. Progressive shared control for training in virtual environments. In *3rd Worldhaptics Conf (WHC)*, pages 332–337. IEEE Press, 2009. → pages 52

[137] C. Liao, F. Guimbretière, and C. E. Loeckenhoff. Pen-top feedback for paper-based interfaces. In *Proceedings of the 19th Annual ACM Symp on User Interface Software and Technology*, UIST, page 201–210, 2006. → pages 31

[138] C. Liao, F. Guimbretière, K. Hinckley, and J. Hollan. Papiercraft: A gesture-based command system for interactive paper. *ACM Trans. Comput.-Hum. Interact.*, (4), 2008. → pages 31

[139] R. C. Limited. Pyqt. 2019. URL https://riverbankcomputing.com/software/pyqt/intro. → pages 59

[140] L.-F. Lin, S.-Y. Teng, R.-H. Liang, and B.-Y. Chen. Stylus assistant: Designing dynamic constraints for facilitating stylus inputs on portable displays. In *SIGGRAPH ASIA 2016 Emerging Technologies*, pages 1–14, 2016. → pages 52, 128

[141] G. L. Long and C. L. Collins. A pantograph linkage parallel platform master hand controller for force-reflection. In *IEEE Int'l Conf on Robotics and Automation (ICRA)*, pages 390–395, 1992. → pages 26

[142] P. Lopes, D. Yuksel, F. Guimbretiere, and P. Baudisch. Muscle-plotter: an interactive system based on electrical muscle stimulation that produces spatial output. In *29th Annual Symp on User Interface Software and Technology*, pages 207–217, 2016. → pages 51, 128, 141

[143] F. M, L. DeLuke, R. Buerba, R. Fan, Y. Zheng, M. Leslie, M. Baumgaertne, and J. Grauer. Haptic biofeedback for improving compliance with lower-extremity partial weight bearing. *Orthopedics*, page 993–998, 2014. → pages 135

[144] K. MacLean and V. Hayward. Do it yourself haptics, part ii: Interaction design. *IEEE Robotics and Automation Society Magazine*, page 104–119, 2008. → pages 74

[145] K. E. MacLean. Haptic interaction design for everyday interfaces. *Human Factors and Ergonomics*, 4(1):149–194, 2008. → pages 28

[146] K. E. MacLean, M. J. Shaver, and D. K. Pai. Handheld haptics: A usb media controller with force sensing. In *IEEE Haptic Symp (HAPTICS)*, pages 311–318, 2002. → pages 20

[147] K. E. MacLean, O. Schneider, and H. Seifi. Multisensory haptic interactions: Understanding the sense and designing for it. In *Handbook of Multimodal-Multisensor Interfaces*, pages 97–142. 2017. → pages 75

[148] A. Magana and S. Balachandran. Unpacking students' conceptualizations through haptic feedback. *Journal of Computer Assisted Learning*, pages 513–531, 2017. → pages 84, 85, 119, 128

[149] M. O. Martinez, T. Morimoto, A. Taylor, A. Barron, J. Pultorak, J. Wang, A. Calasanz-Kaiser, R. Davis, P. Blikstein, and A. Okamura. 3d printed haptic devices for educational applications. In *IEEE Haptics Symp (HAPTICS)*, pages 126–133, 2016. → pages 20, 28, 44, 74

[150] T. H. Massie and J. K. Salisbury. The phantom haptic interface: A device for probing virtual objects. In *Proceedings of the ASME Dynamic Systems and Control Division*, pages 295–301, 1994. → pages 16, 18

[151] A. Meier, H. Spada, and N. Rummel. A rating scheme for assessing the quality of computer-supported collaboration processes. *International Journal of Computer-Supported Collaborative Learning*, pages 63–86, 2007. → pages xvi, 76, 99, 102

[152] T. Mikropoulos and I. Bellou. Educational robotics as mindtools. *Themes in Science and Technology Education*, pages 5–14, 2013. ISSN 1792-8788. → pages 1

[153] B. E. Miller, J. E. Colgate, and R. A. Freeman. On the role of dissipation in haptic systems. *IEEE Trans on Robotics*, page 768–771, 2004. → pages 143

[154] G. Minaker, O. Schneider, R. Davis, and K. E. MacLean. Handson: enabling embodied, creative stem e-learning with programming-free force feedback. In *Eurohaptics*, pages 427–437, London, 2016. → pages 123, 129

[155] G. Minaker, O. Schneider, R. Davis, and K. E. MacLean. Handson: Enabling embodied, creative stem e-learning with programming-free force feedback. In *Eurohaptics*, pages 427–437. Springer, 2016. → pages 28, 74

[156] J. Minogue and G. Jones. Measuring the impact of haptic feedback using the solo taxonomy. *International Journal of Science Education*, pages 1359–1378, 2009. → pages 74, 114, 128

[157] J. Minogue, M. G. Jones, B. Broadwell, and T. Oppewall. The impact of haptic augmentation on middle school students' conceptions of the animal cell. *Virtual Reality*, pages 293–305, 2006. → pages 74

[158] M. Minsky, M. Ouh-young, O. Steele, F. P. B. Jr., and M. Behensky. Feeling and seeing: issues in force display. *ACM Computer Graphics*, page 235–242, 1990. → pages 17

[159] M. D. R. Minsky. *Computational haptics: the sandpaper system for synthesizing texture for a force-feedback display*. PhD thesis, Queen's University, 1995. → pages 34

[160] G. J. Monkman. An electrorheological tactile display. *Presence: Teleoperators and Virtual Environments*, page 219–228, 1992. → pages 25

[161] M. Montessori and B. Carter. *The secret of childhood*. Orient Longmans Calcutta, 1936. → pages 123

[162] O. Mubin, C. J. Stevens, S. Shahid, A. Al Mahmud, and J. Dong. A review of the applicability of robots in education. *Technology for Education and Learning*, 1:1–7, 2013. → pages 2

[163] P. A. Mueller and D. M. Oppenheimer. The pen is mightier than the keyboard: Advantages of longhand over laptop note taking. *Psychological Science*, page 1159–1168, 2014. → pages 47

[164] J. Mullins, C. Mawson, and S. Nahavandi. Haptic handwriting aid for training and rehabilitation. In *IEEE Int'l Conf on Systems, Man and Cybernetics*, pages 2690–2694, 2005. → pages 23, 26, 128

[165] J. Murayama, L. Bougrila, Y. Luo, K. Akahane, S. Hasegawa, B. Hirsbrunner, and M. Sato. Spidar: a two-handed haptic interface for bimanual vr interaction. In *Proceedings of EuroHaptics*, volume 2004, pages 138–146. Citeseer, 2004. → pages 128

[166] K. Nakagaki and Y. Kakehi. Comp*pass: A compass-based drawing interface. In *CHI Extended Abstracts on Human Factors in Computing Systems*, CHI, pages 447–450, 2014. → pages 51, 128

[167] A. Nomoto, Y. Ban, T. Narumi, T. Tanikawa, and M. Hirose. Supporting precise manual-handling task using visuo-haptic interaction. In *Proceedings of the 7th Augmented Human International Conf*, pages 1–10. ACM, 2016. → pages 51

[168] B. M. of Education". "areas of learning: Science – science for citizens 11", 2018. URL https://curriculum.gov.bc.ca/sites/curriculum.gov.bc.ca/files/curriculum/science/en_science_11_science-for-citizens_elab.pdf. → pages 82

[169] M. A. Otaduy and M. C. Lin. High fidelity haptic rendering. *Synthesis Lectures on Computer Graphics*, pages 1–112, 2006. → pages 16, 17

[170] A. Özgür, W. Johal, F. Mondada, and P. Dillenbourg. Windfield: Learning wind meteorology with handheld haptic robots. In *Proceedings of the 2017 ACM/IEEE International Conf on Human-Robot Interaction*, HRI, page 156–165, 2017. → pages 7, 97

[171] A. Özgür, W. Johal, F. Mondada, and P. Dillenbourg. Haptic-enabled handheld mobile robots: Design and analysis. In *Proceedings of the CHI Conf on Human Factors in Computing Systems*, CHI, pages 2449–2461, 2017. → pages xix, 6, 27, 37, 76, 80, 86, 121

[172] A. Özgür, S. Lemaignan, W. Johal, M. Beltran, M. Briod, L. Pereyre, F. Mondada, and P. Dillenbourg. Cellulo: Versatile handheld robots for education. In *ACM/IEEE Int'l Conf on Human-Robot Interaction (HRI)*, pages 119–127, 2017. → pages 124, 128

[173] P. Paczkowski, M. H. Kim, Y. Morvan, J. Dorsey, H. Rushmeier, and C. O'Sullivan. Insitu: Sketching architectural designs in context. *ACM Trans. Graph.*, page 1–10, 2011. → pages 30

[174] S. Papert. *Mindstorms: Children, Computers, and Powerful Ideas*. Basic Books, New York, 1980. → pages 1, 2, 6, 70, 106, 123, 126

[175] S. Papert and I. Harel. Situating constructionism. In S. Papert and I. Harel, editors, *Constructionism*, chapter 1, page 3. Ablex Publishing Corporation, 1991. → pages 123

[176] G. Park, H. Cha, and S. Choi. Attachable and detachable vibrotactile feedback modules and their information capacity for spatiotemporal patterns. In *IEEE WorldHaptics*, pages 78–83, Munich, 2017. → pages 28

[177] H. M. Paynter. *Analysis and design of engineering systems*. MIT Press, 1960. → pages 150

[178] K. Perkins, W. Adams, M. Dubson, N. Finkelstein, S. Reid, C. Wieman, and R. LeMaster. Phet: Interactive simulations for teaching and learning physics. *The physics teacher*, pages 18–23, 2006. → pages 18

[179] K. Perlin, Z. He, and K. Rosenberg. Chalktalk: A visualization and communication language–as a tool in the domain of computer science education. *arXiv preprint arXiv:1809.07166*, 2018. → pages 30

[180] J. Piaget. Part i: Cognitive development in children: Piaget development and learning. *Journal of research in science teaching*, pages 176–186, 1964. → pages 6, 123, 152

[181] J. Piaget. *The development of thought: Equilibration of cognitive structures.(Trans A. Rosin).* Viking, 1977. → pages 123

[182] M. Porcheron, J. E. Fischer, and S. Sharples. Using mobile phones in pub talk. In *Proceedings of the 19th ACM Conf on Computer-Supported Cooperative Work amp; Social Computing*, page 1649–1661, 2016. → pages 89

[183] M. Price and F. C. Sup. A robotic touchscreen totem for two-dimensional haptic force display. In *2016 IEEE Haptics Symp (HAPTICS)*, pages 72–77, 2016. → pages 21

[184] I. Radu and B. Schneider. What can we learn from augmented reality (ar)? In *Proceedings of the 2019 CHI Conf on Human Factors in Computing Systems*, pages 1–12, 2019. → pages 84, 112

[185] I. Radu, V. Hv, and B. Schneider. Unequal impacts of augmented reality on learning and collaboration during robot programming with peers. pages 1–23, 2021. → pages 119, 120

[186] C. Ramstein and V. Hayward. The pantograph: a large workspace haptic device for a multi-modal human-computer interaction. In *ACM Conf on Human Factors in Computing Systems (CHI)*, Boston, MA, 1994. → pages 20, 23

[187] M. Reiner. Sensory cues, visualization and physics learning. *International Journal of Science Education*, 31(3):343–364, 2009. → pages 126

[188] M. D. Renken and N. Nunez. Computer simulations and clear observations do not guarantee conceptual understanding. *Learning and Instruction*, pages 10–23, 2013. → pages 128

[189] M. Resnick. Technologies for lifelong kindergarten. *Educational technology research and development*, pages 43–55, 1998. → pages 122

[190] M. Resnick, F. Martin, R. Berg, R. Borovoy, V. Colella, K. Kramer, and B. Silverman. Digital manipulatives: New toys to think with. pages 281–287. → pages 70, 124, 126, 129

[191] M. J. Rodríguez-Triana, L. P. Prieto, A. Vozniuk, M. S. Boroujeni, B. A. Schwendimann, A. Holzer, and D. Gillet. Monitoring, awareness and reflection in blended technology enhanced learning: a systematic review. *International Journal of Technology Enhanced Learning*, 9(2-3):126–150, 2017. → pages 126

[192] H. Romat, N. Henry Riche, K. Hinckley, B. Lee, C. Appert, E. Pietriga, and C. Collins. Activeink: (th)inking with data. In *Proceedings of the CHI Conf on Human Factors in Computing Systems*, CHI, page 1–13, 2019. → pages 121

[193] J. F. Roscoe, S. Fearn, and E. Posey. Teaching computational thinking by playing games and building robots. In *2014 International Conf on Interactive Technologies and Games*, pages 9–12, 2014. → pages 3

[194] L. B. Rosenberg. Virtual fixtures: Perceptual tools for telerobotic manipulation. In *Virtual Reality Annual International Symp*, pages 76–82. IEEE, 1993. → pages 17

[195] L. B. Rosenberg. Haptic feedback stylus and other devices, May 3 2007. URL https://patents.google.com/patent/US7508382B2/en?oq=US+Pat.+No+7%2c508%2c382+B2. US Patent US7265750B2. → pages 22

[196] J. Saldien, K. Goris, S. Yilmazyildiz, W. Verhelst, and D. Lefeber. On the design of the huggable robot probo. *Physical Agents*, 2(2):3–11, 2008. → pages 72

[197] J. K. Salisbury and M. A. Srinivasan. Phantom-based haptic interaction with virtual objects. *Computer Graphics and Applications, IEEE*, page 6–10, 1997. → pages 137

[198] J. K. Salisbury, F. Conti, and F. Barbagli. Haptic rendering: introductory concepts. *IEEE Computer Graphics & Applications*, 24(2):24–32, 2004. ISSN 0272-1716. doi:10.1109/MCG.2004.1274058. → pages 35

[199] B. Schneider, P. Jermann, G. Zufferey, and P. Dillenbourg. Benefits of a tangible interface for collaborative learning and interaction. *IEEE Trans on Learning Technologies*, pages 222–232, 2011. → pages 77

[200] Y. Sefidgar, K. E. MacLean, S. Yohanan, M. Van der Loos, E. A. Croft, and J. Garland. Design and evaluation of a touch-centered calming interaction with a social robot. *Trans on Affective Computing*, page 108–121, 2015. → pages 72

[201] H. Seifi, G. Park, F. Fazlollahi, J. Sastrillo, J. Ip, A. Agrawal, K. Kuchenbecher, and K. E. MacLean. Haptipedia: Accelerating haptic device discovery to support interaction engineering design. In *Conf on Human Factors in Computing Systems (CHI)*, pages 12–24. ACM, 2019. → pages 19, 47, 74, 75

[202] U. A. Shaikh, A. J. Magana, L. Neri, D. Escobar-Castillejos, J. Noguez, and B. Benes. Undergraduate students' conceptual interpretation and perceptions of haptic-enabled learning experiences. *International Journal of Educational Technology in Higher Education*, page 15, 2017. → pages 87

[203] Y. shan Chang, M. Y.-C. Chen, M.-J. Chuang, and C. hui Chou. Improving creative self-efficacy and performance through computer-aided design application. *Thinking Skills and Creativity*, pages 103 – 111, 2019. → pages 49

[204] C. Shannon. Communication in the presence of noise. *Proceedings of the IRE*, pages 10–21, jan 1949. → pages 137

[205] G. Sidorov, A. Gelbukh, H. Gómez-Adorno, and D. Pinto. Soft similarity and soft cosine measure: Similarity of features in vector space model. *Computación y Sistemas*, 18(3):491–504, 2014. → pages 97

[206] B. Smits-Engelsman, A. Niemeijer, and G. van Galen. Fine motor deficiencies in children diagnosed as dcd based on poor grapho-motor ability. *Human Movement Science*, pages 161–182, 2001. → pages 135

[207] S. Snibbe, S. Anderson, and B. Verplank. Springs and constraints for 3d drawing. In *Third Phantom User's Group*, volume M.I.T. Artificial Intelligence Laboratory Technical Report AITR-1643, pages 1–4, 1998. → pages 49, 52, 56

[208] H. Song, T. Grossman, G. Fitzmaurice, F. Guimbretiere, A. Khan, R. Attar, and G. Kurtenbach. Penlight: Combining a mobile projector and a digital pen for dynamic visual overlay. In *Proceedings of the SIGCHI Conf on Human Factors in Computing Systems*, CHI, pages 143–152, 2009. → pages 51

[209] L. J. Speed and A. Majid. Grounding language in the neglected senses of touch, taste, and smell. *Cognitive neuropsychology*, pages 363–392, 2020. → pages 74

[210] M. A. Srinivasan. *Haptic Interfaces*. Washington, D.C.: Report of the Committee on Virtual Reality Research and Development, National Research Council, National Academy Press, 1995. → pages 16

[211] J. Steimle. Designing pen-and-paper user interfaces for interaction with documents. In *Proceedings of the 3rd International Conf on Tangible and Embedded Interaction (TEI)*, pages 197–204. ACM Press, 2009. → pages 52, 53

[212] I. E. Sutherland. Sketchpad a man-machine graphical communication system. *Simulation*, 2:3–20, 1964. → pages 28

[213] H. Z. Tan and A. Pentland. Tactual displays for wearable computing. In *International Symp on Wearable Computers*, pages 84–89. IEEE Computer Society, 1997. → pages 20

[214] H. Z. Tan, M. A. Srinivasan, B. Eberman, and B. Cheng. Human factors for the design of force-reflecting haptic interfaces. In *3rd Ann. Symp. on Haptic Interfaces for Virtual Environment and Teleoperator Systems, ASME/IMECE*, pages 353–359, 1994. → pages 17

[215] A. Teranishi, G. Korres, W. Park, and M. Eid. Combining full and partial haptic guidance improves handwriting skills development. *IEEE Trans on Haptics*, pages 509–517, 2018. → pages 52, 135

[216] D. Tsetserukou, A. Neviarouskaya, H. Prendinger, N. Kawakami, and S. Tachi. Affective haptics in emotional communication. In *2009 3rd International Conf on Affective Computing and Intelligent Interaction and Workshops*, pages 1–6, 2009. → pages 72

[217] D. Tsetserukou, K. Sato, and S. Tachi. Exointerfaces: novel exosceleton haptic interfaces for virtual reality, augmented sport and rehabilitation. In *Proceedings of the 1st Augmented Human International Conf*, pages 1–6, 2010. → pages 128

[218] D. G. Ullman, S. Wood, and D. Craig. The importance of drawing in the mechanical design process. *Computers Graphics*, pages 263 – 274, 1990. → pages 30

[219] P. Vandoren, T. Van Laerhoven, L. Claesen, J. Taelman, C. Raymaekers, and F. Van Reeth. Intupaint: Bridging the gap between physical and digital painting. In *2008 3rd IEEE International Workshop on Horizontal Interactive Human Computer Systems*, pages 65–72, 2008. → pages 47

[220] D. Vaquero-Melchor and A. M. Bernardos. Enhancing interaction with augmented reality through mid-air haptic feedback: architecture design and user feedback. *Applied Sciences*, page 5123, 2019. → pages 128

[221] S. A. Wall and S. A. Brewster. *Assessing Haptic Properties for Data Representation*, page 858–859. 2003. → pages 128

[222] D. Wang, Y. Zhang, and C. Yao. Stroke-based modeling and haptic skill display for chinese calligraphy simulation system. *Virtual Reality*, pages 118–132, 2006. → pages 128

[223] M. Weiser. The computer for the 21st century. *Scientific American*, page 94–110, 1991. → pages 151, 153

[224] M. Weiss, C. Wacharamanotham, S. Voelker, and J. Borchers. Fingerflux: near-surface haptic feedback on tabletops. *Proceedings of the 24th annual ACM on User interface software and technology*, page 615–620, 2011. → pages 24

[225] E. N. Wiebe, J. Minogue, M. Gail Jones, J. Cowley, and D. Krebs. Haptic feedback and students' learning about levers: Unraveling the effect of simulated touch. *Computers Education*, pages 667 – 676, 2009. → pages 78

[226] Wikipedia. The blind men and the elephant (buddhist parable). 2021, 2021. URL https://en.wikipedia.org/wiki/Blind_men_and_an_elephant. → pages 74

[227] L. Winfield, J. Glassmire, J. E. Colgate, and M. Peshkin. T-pad: Tactile pattern display through variable friction reduction. In *Second Joint EuroHaptics Conf and Symp on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC)*, pages 421–426, 2007. → pages 25

[228] W. Witaya, A. Prasad, P. Michael, and J. E. Colgate. Cobots: A novel material handling technology. In *7th Ann. Symp. on Haptic Interfaces for Virtual Environment and Teleoperator Systems, ASME/IMECE*, Anaheim, USA, 1998. → pages 21

[229] E. B. Witherspoon, R. M. Higashi, C. D. Schunn, E. C. Baehr, and R. Shoop. Developing computational thinking through a virtual robotics programming curriculum. *ACM Trans. Comput. Educ.*, pages 1–4, 2017. → pages 3

[230] J. O. Wobbrock, A. D. Wilson, and Y. Li. Gestures without libraries, toolkits or training: A $1 recognizer for user interface prototypes. In *Proceedings of the 20th Annual ACM Symp on User Interface Software and Technology*, UIST, page 159–168, 2007. → pages 139

[231] J. Yamaoka and Y. Kakehi. depend: Augmented handwriting system using ferromagnetism of a ballpoint pen. In *26th Annual ACM Symp on User Interface Software and Technology*, pages 203–210. ACM, 2013. → pages 26, 51, 128

[232] S. Yohanan and K. E. MacLean. The role of affective touch in human-robot interaction: Human intent and expectations in touching the haptic creature. *International Journal of Social Robotics*, pages 163–180, 2012. → pages 72, 89

[233] Younq-Seok, M. Collins, W. Bulmer, S. Sharma, and J. Mayrose. Haptics assisted training (hat) system for children's handwriting. In *2013 World Haptics Conf (WHC)*, pages 559–564, 2013. → pages 128

[234] Z. C. Zacharia. Examining whether touch sensory feedback is necessary for science learning through experimentation: A literature review of two different lines of research across k-16. *Educational Research Review*, pages 116 – 137, 2015. → pages 77, 78

[235] Z. C. Zacharia and M. Michael. Using physical and virtual manipulatives to improve primary school students' understanding of concepts of electric circuits. In *New developments in science and technology education*, pages 125–140. Springer, 2016. → pages 128

[236] Z. C. Zacharia, E. Loizou, and M. Papaevripidou. Is physicality an important aspect of learning through science experimentation among kindergarten students? *Early Childhood Research Quarterly*, pages 447–457, 2012. → pages 78, 119, 128

[237] R. Zeleznik, A. Bragdon, F. Adeputra, and H.-S. Ko. Hands-on math: A page-based multi-touch and pen desktop for technical work and problem solving. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology*, pages 17–26, 2010. → pages 31

[238] K. Zhao and C. K. Chan. Fostering collective and individual learning through knowledge building. *International Journal of Computer-Supported Collaborative Learning*, pages 63–95, 2014. → pages 78

[239] M. Zinn, O. Khatib, B. Roth, and J. K. Salisbury. Large workspace haptic devices - a new actuation approach. *Symp on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 185–192, 2008. → pages 17

[240] R. F. Zory, D. Boërio, M. Jubeau, and N. A. Maffiuletti. Central and peripheral fatigue of the knee extensor muscles induced by electromyostimulation. *International journal of sports medicine*, pages 847–853, 2005. → pages 51

Appendix A

MagicPen Technical Demonstration

Here, we present additional data on the technical performance of the MagicPen's stylus form factor (Figure 3.1(a)) presented in Chapter 3.

The graphs here showcase a user interaction with different virtual objects via MagicPen. We hope that these graphs present both temporal and spatial aspects of haptic rendering with our device.

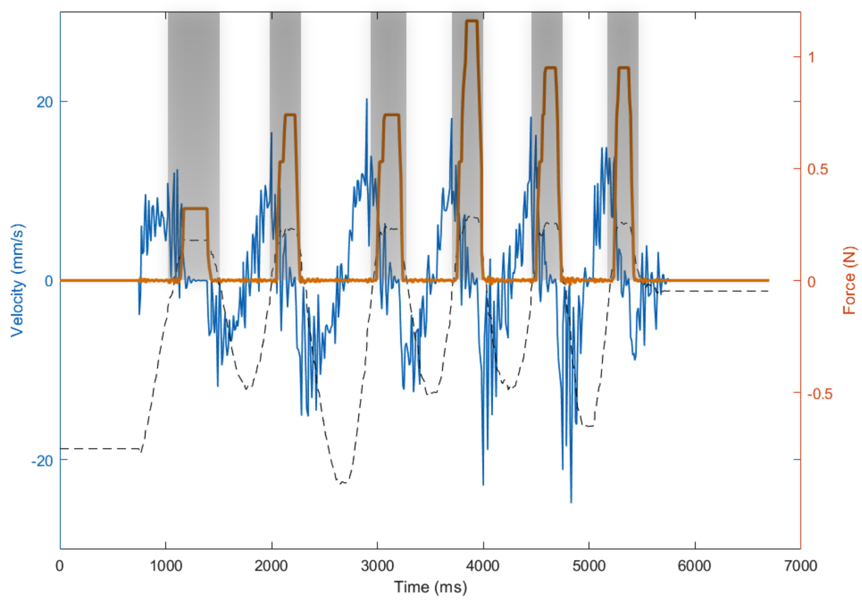

Figure A.1: The force and velocity behaviour as the user runs the avatar into a virtual wall. The blue line represents the velocity, the orange line shows the force, and the black line represents the user's trajectory. The gray box represents the regions where the user's avatar is inside the virtual wall.

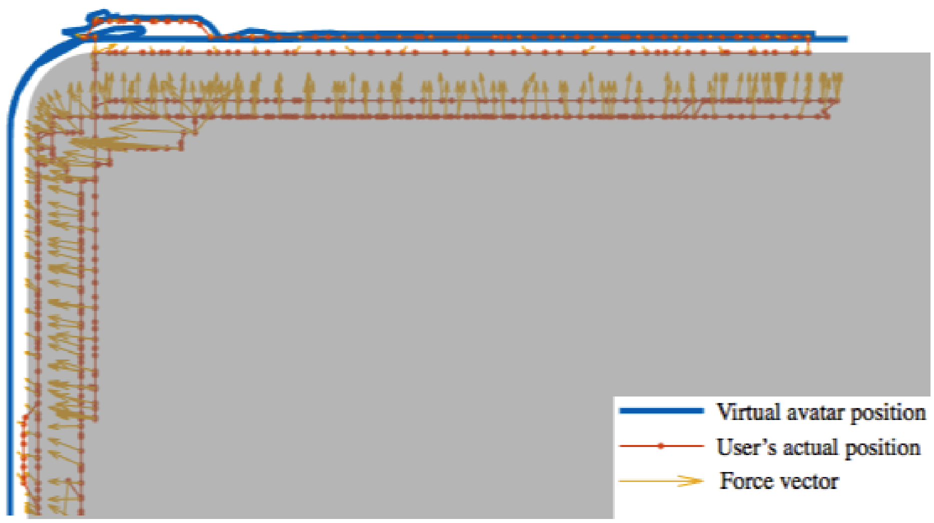

Figure A.2: A user is moving the MagicPen around the corner of a box (gray).

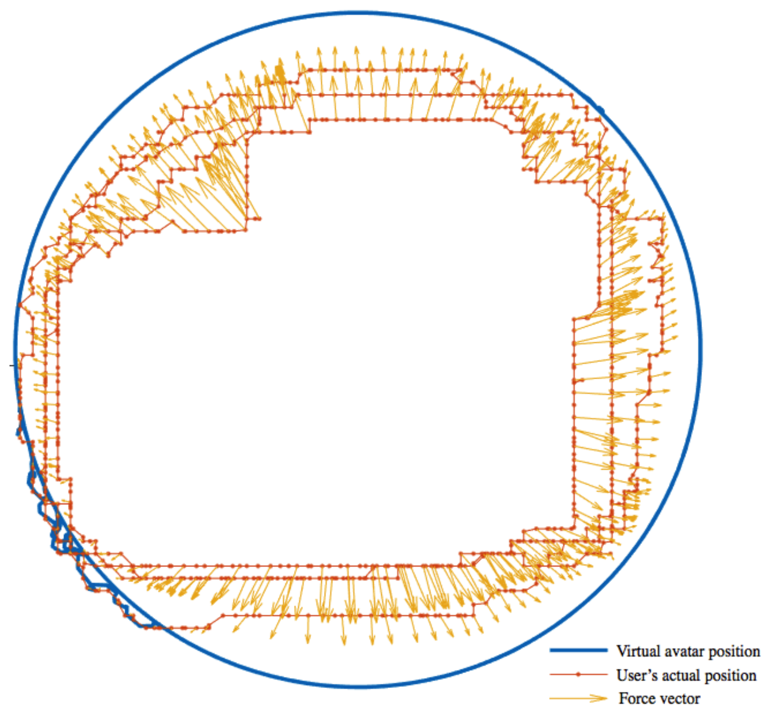

Figure A.3: A user is moving the MagicPen around a circular wall.

Appendix B

Phasking Experiment Results

The drawings here presents the results of drawing experiments. *P1-P7 are novices, and E1-E3 are domain experts.*

- Neo: No assist (NeoSmart pen)
- Bring: Physical assists-active constraint
- Bound: physical assists-passive constraint
- SC: shared control
- NSC: No shared controla

Activities:

1. Draw bring constraints : a straight line (see Section B.1)
2. Draw bring constraints : a rectangle (see Section B.2)
3. Draw bring constraints : a rectangle in perspective (see Section B.3)
4. Draw bring constraints : a circle (see Section B.4)
5. Bound constraints : lines meeting an invisible line barrier (see Section B.5)
6. Shared control : a sine wave/invisible line barrier at the center (see Section B.6)

B.1 Draw bring constraints : a straight line

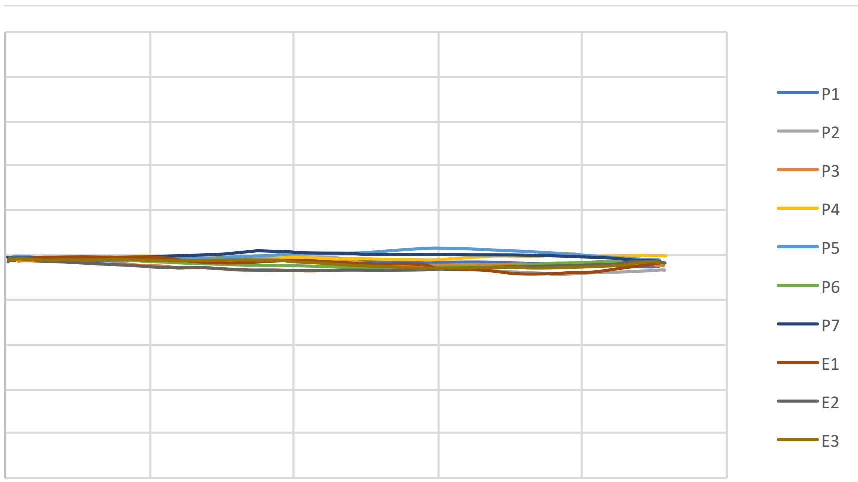

Figure B.1: Line with NeoSmart pen.

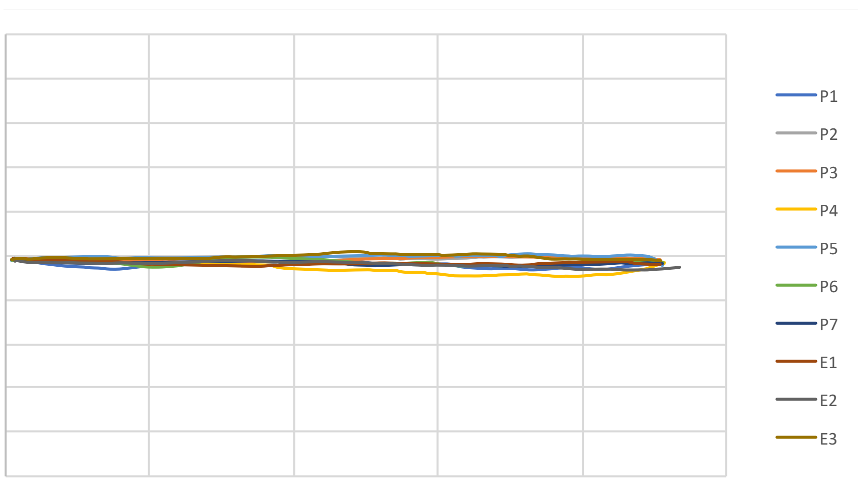

Figure B.2: Line – bring.

B.2 Draw bring constraints : a rectangle

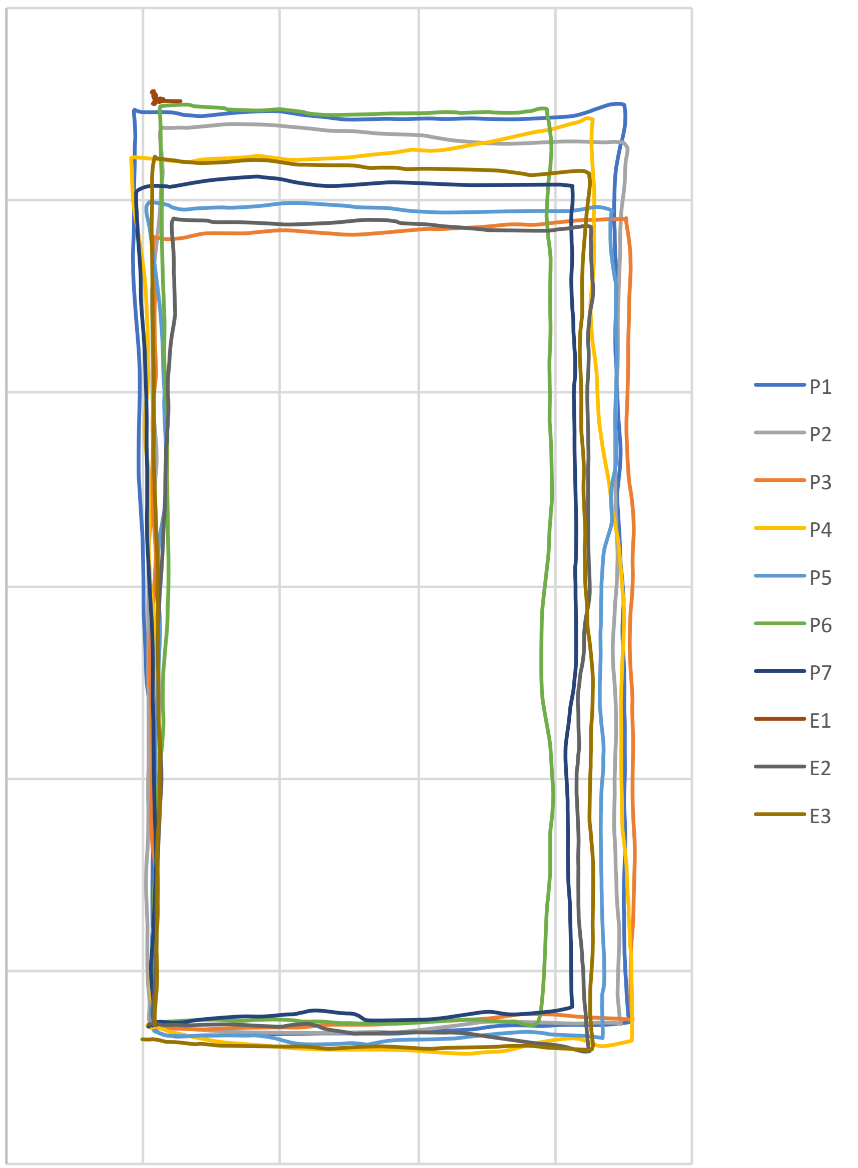

Figure B.3: Rectangle with NeoPen

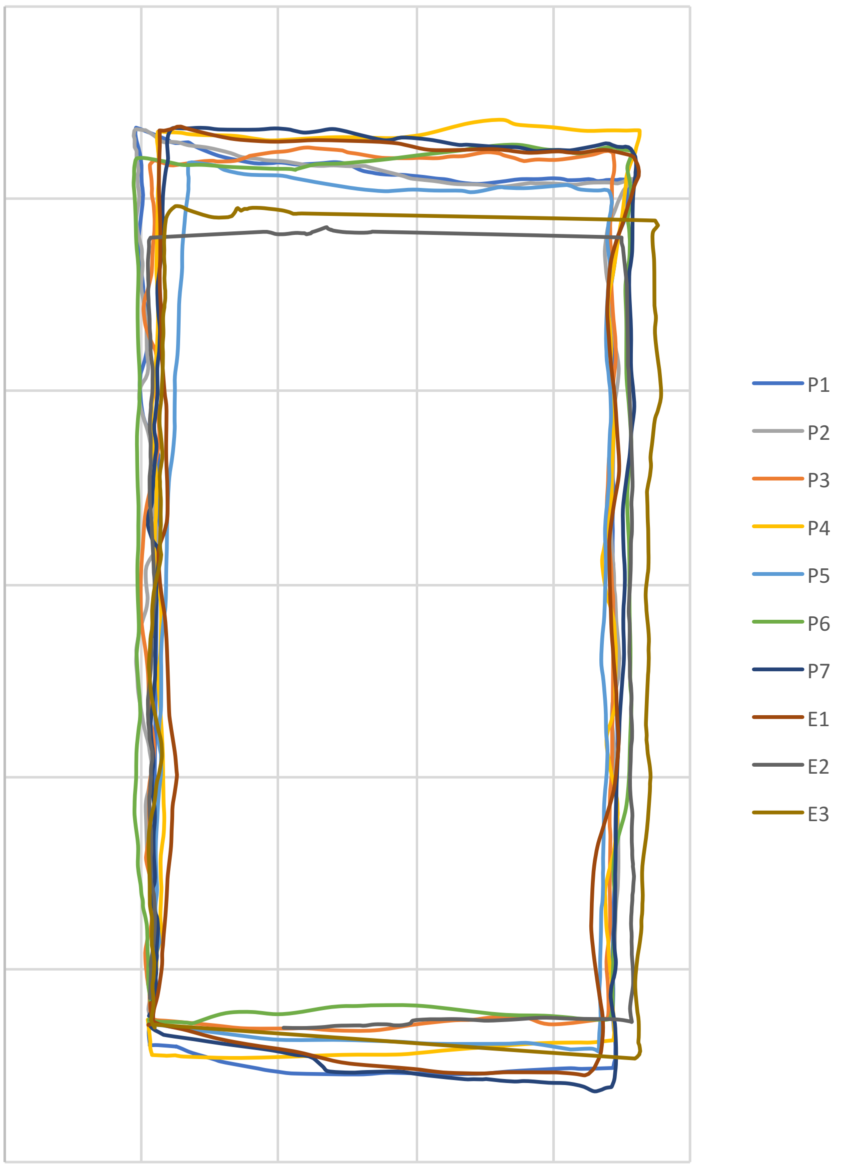

Figure B.4: Rectangle – bring

B.3 Draw bring constraints : a rectangle in perspective

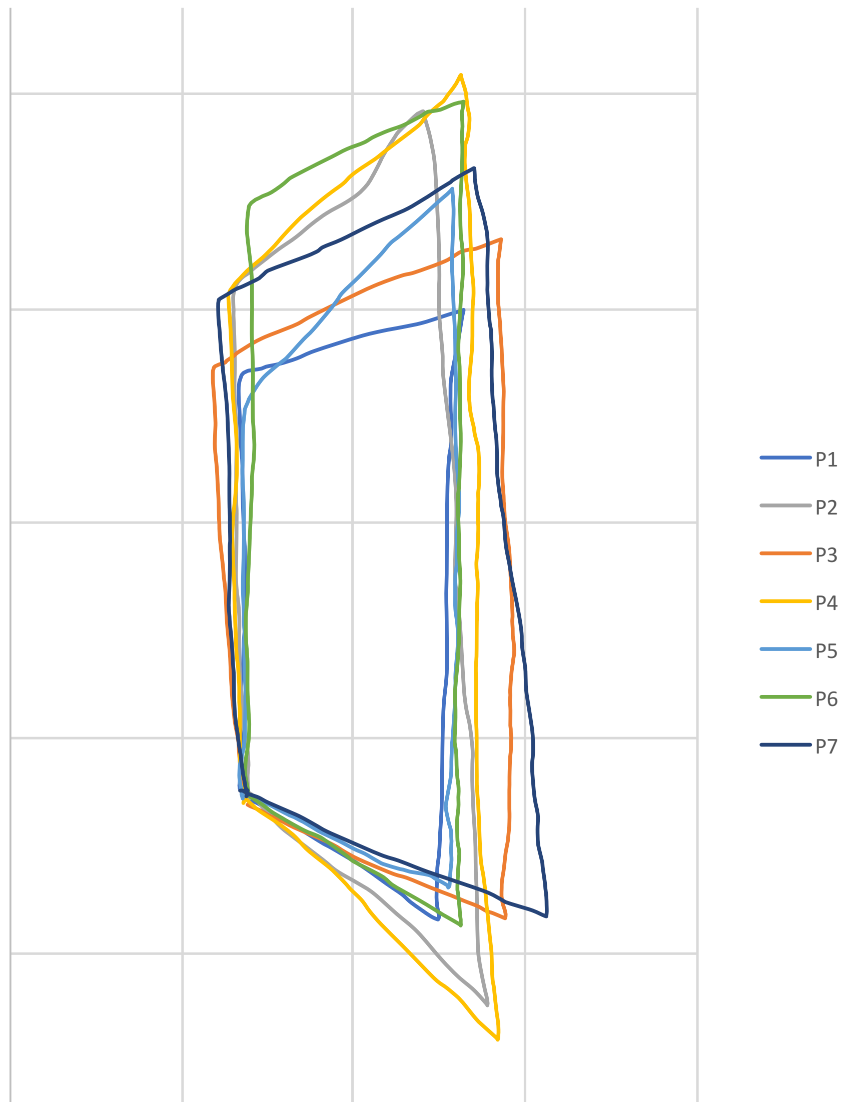

Figure B.5: Rectangle in perspective with NeoSmart pen

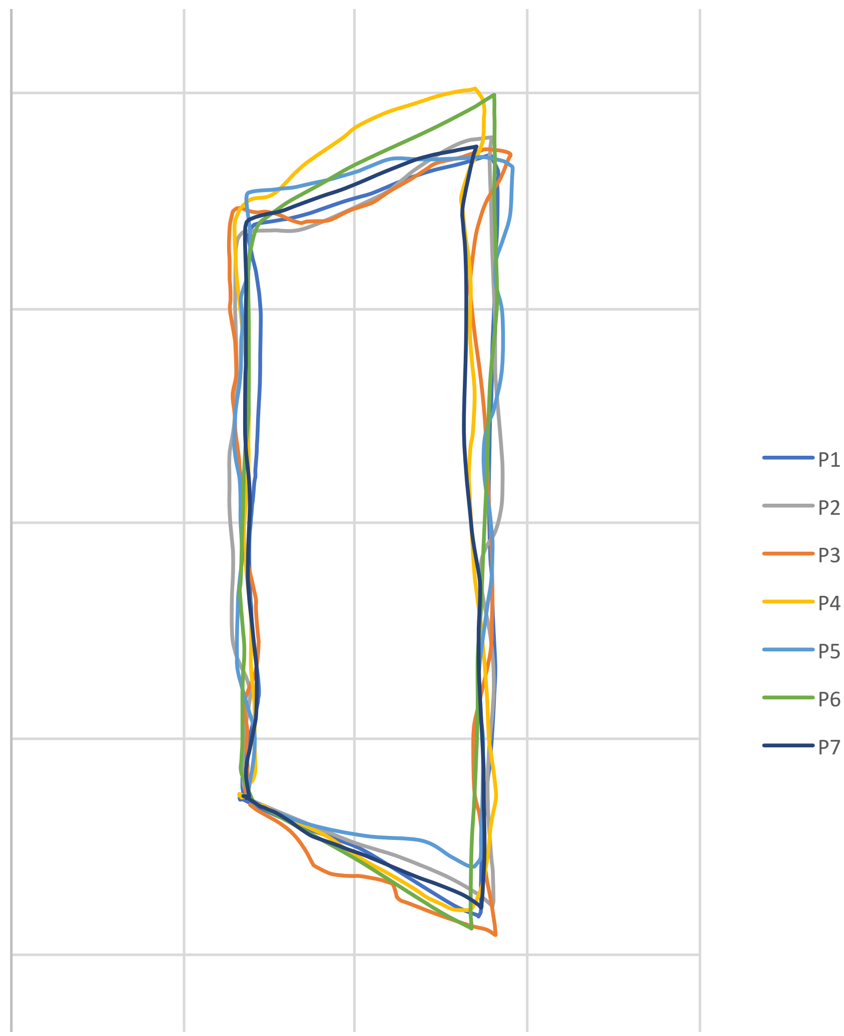

Figure B.6: Rectangle in perspective – bring

B.4 Draw bring constraints : a circle

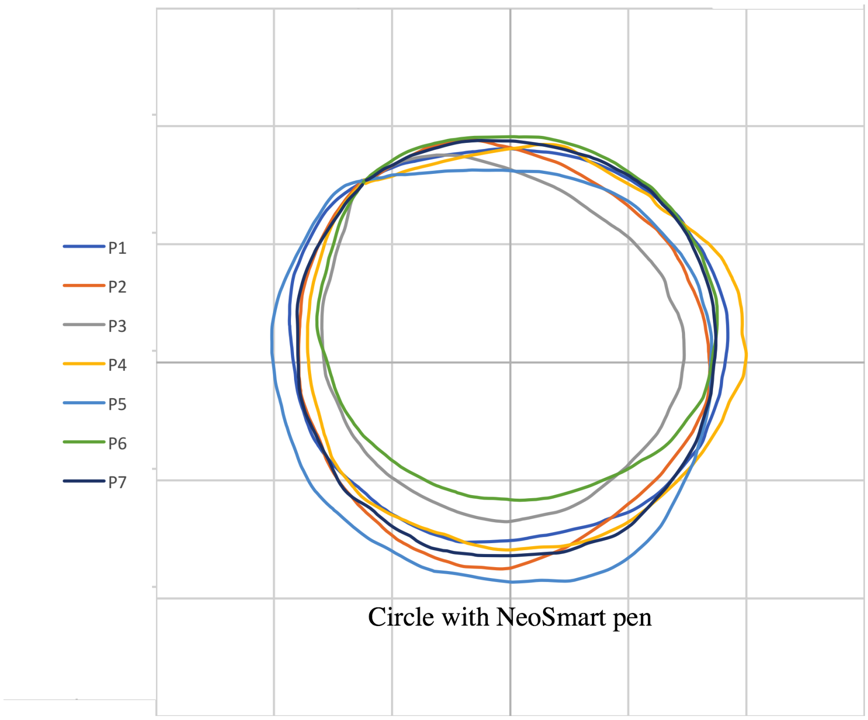

Figure B.7: Circle with NeoSmart pen

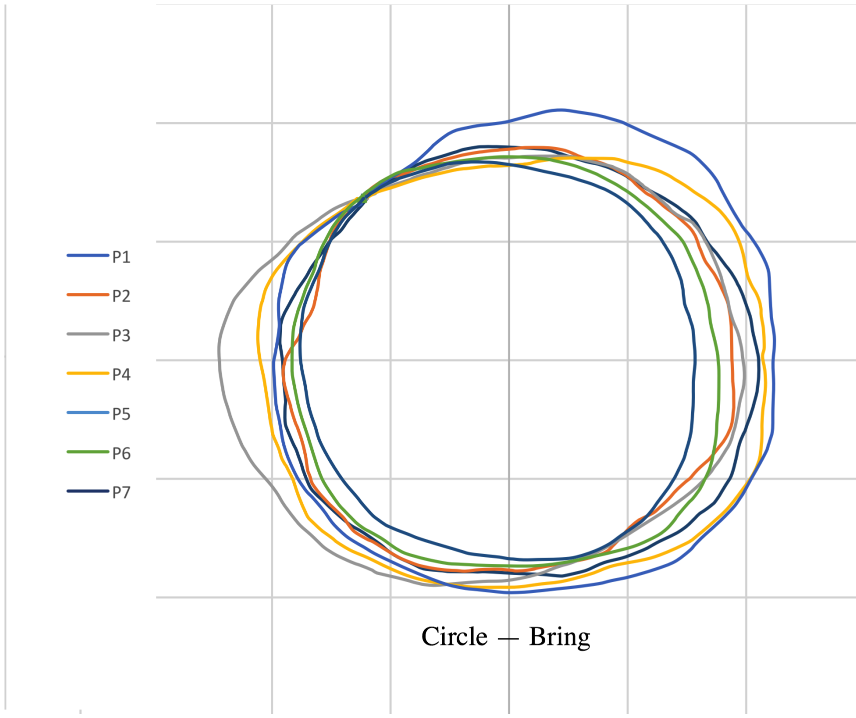

Figure B.8: Circle – bring

B.5 Bound constraints:lines meeting an invisible line barrier

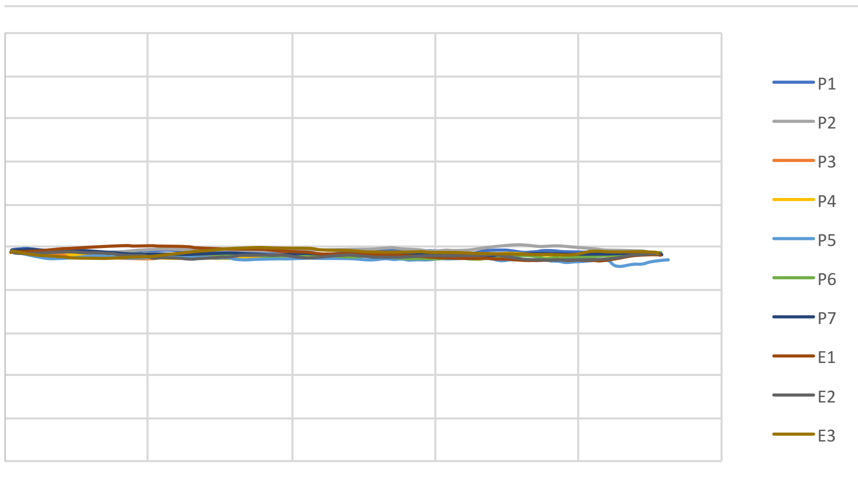

Figure B.9: Line – bound

B.6 Shared control : a sine wave/invisible line barrier at the center

()

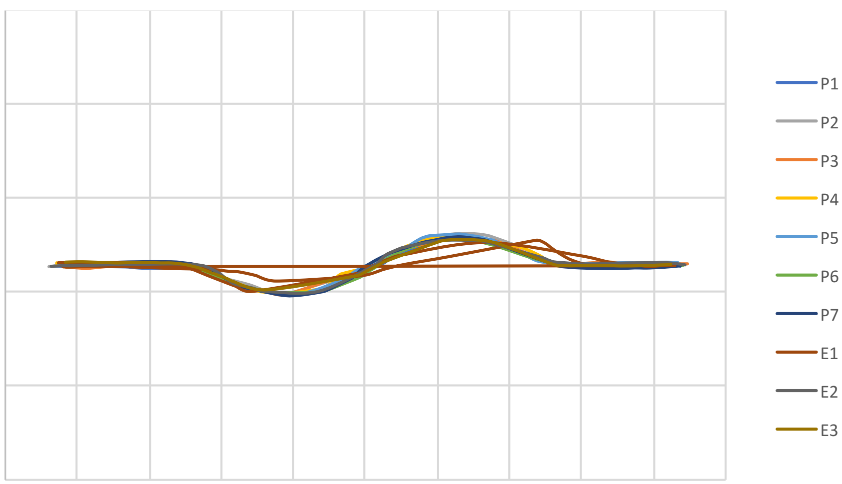

Figure B.10: A sine wave with no shared control – bound

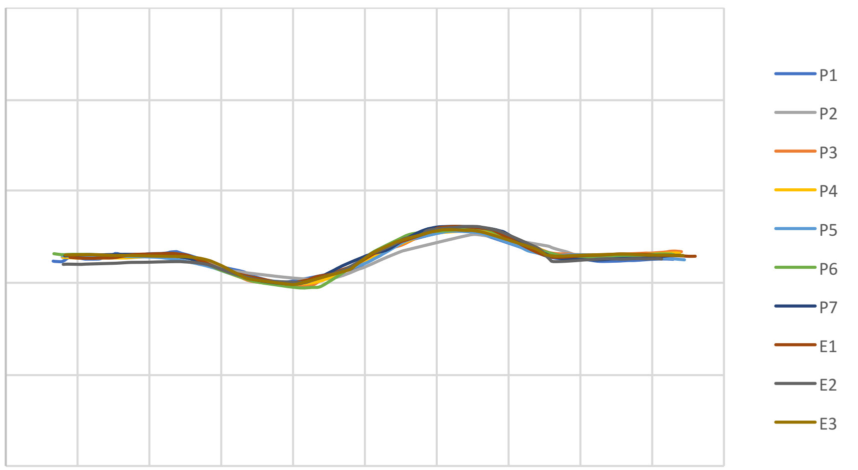

Figure B.11: A sine wave with shared control – bound

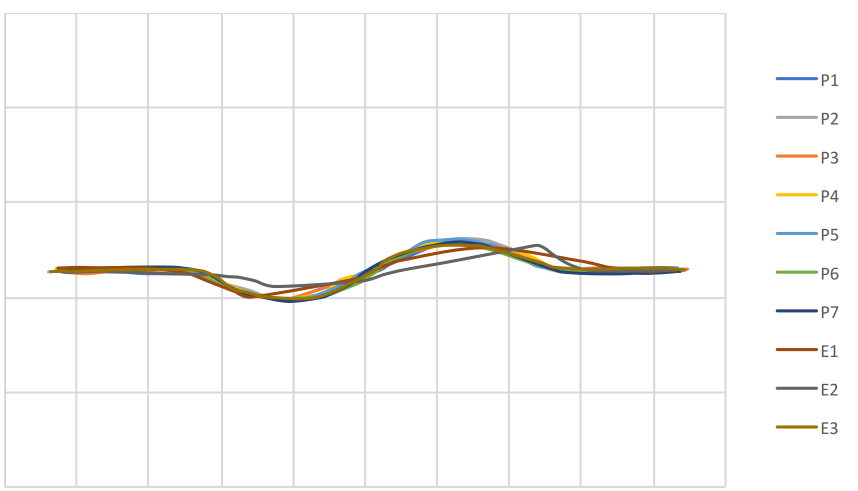

Figure B.12: A sine wave with no shared control – bring

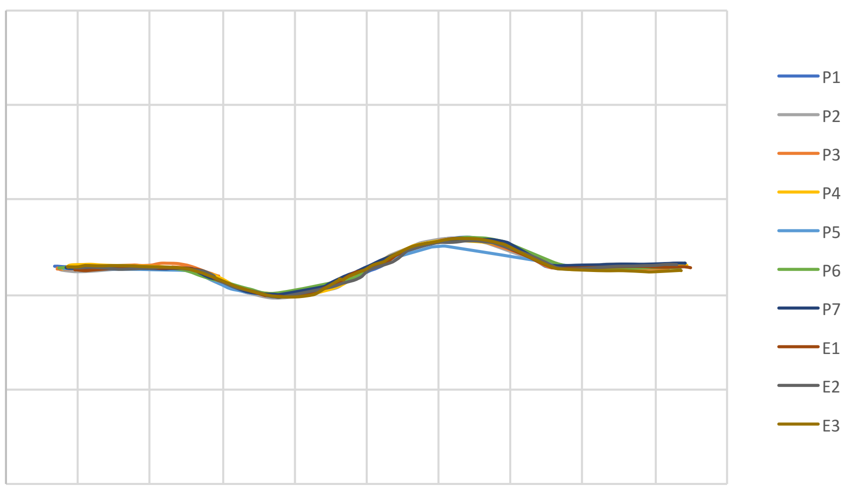

Figure B.13: A sine wave with shared control – bring

Appendix C

Collaborative Grounding Pre-test/Post-test Questions

Electrostatic Lab Pre-Test

Name :.. Question purpose in yellow (removed in the original version)

Group :..

Please work individually to answer the following questions. There is no grading, so just answer to the best of your knowledge.

Determining existing knowledge about electrostatic forces & their interactions

1. In the context of physics, how would you define the following words?

   Electrostatic force: ____

   Net force: ____

   System equilibrium: ____

Testing understanding of net forces and how to calculate them

2. Three small spheres are lined up in a row as shown in the figure below. A negatively charged sphere is placed in the middle of the two positive charges. The charge on this sphere is -2 microCoulombs.
   Circle the sphere that experiences the lowest net force:

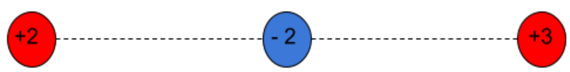

Testing understanding of direction for net forces + equilibrium requirements

3. Three tiny spheres with identical charges of 5 microCoulombs are situated at the corners of an equilateral triangle. Draw the net force acting on each sphere:

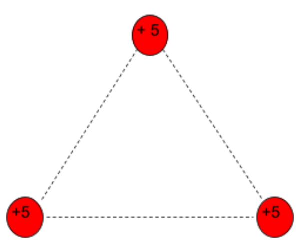

Is the system in an equilibrium state?

- ❏ Yes
- ❏ No
- ❏ I don't know

Figure C.1: Electrostatic lab pre-test

Electrostatic Lab Post-Test

Name :.. Question purpose in yellow (removed in the original version)

Group: ....

Please work individually to answer the following questions. There is no grading, so just answer to the best of your knowledge.

1. Briefly, how would you define the following words, in terms of physics?

   Electrostatic force: ____

   Net force: ____

   System equilibrium: ____
   Comparing definitions of key terms pre- and post-learning activity

2. Three small spheres are lined up in a row as shown in the figure below. The first and the third spheres are 10 cm apart and have the same charge of -8 microCoulombs. A positively charged sphere is placed in the middle of the two negative charges. The charge on this sphere is -8 microCoulombs.
   Draw the force vectors acting on each sphere.
   What is the net force on the positively charged sphere? _______________

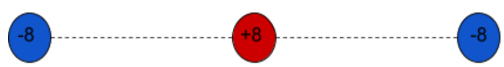

Testing understanding of net force

3. Three tiny spheres with identical charges of +3 microCoulombs are situated at the corners of a square. Is the system below in equilibrium?

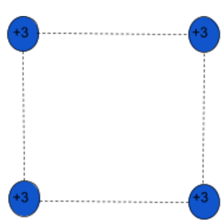

- ❏ Yes
- ❏ No
- ❏ I don't know

Testing understanding of system equilibrium

Figure C.2: Electrostatic lab post-test

Momentum & Collision Lab Pre-Test

Name :.. Question purpose in yellow (removed in the original version)

Group :...

Please work by yourself to answer the following questions:

1. In the context of physics, how would you define the following words?

   Mass: ____

   Velocity: ____

   Momentum: ____

   Determining existing knowledge about momentum, mass and velocity

1. Circle the object with the most momentum:

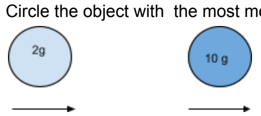

Testing understanding of p = mv

3. Circle the object with the most momentum:

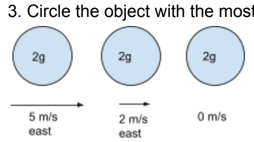

Testing understanding of p = mv

3. Two objects, one travelling east and the other travelling at west, collide. After the collision, they stick together and both travel east. This collision is:
   a. Elastic
   b. Inelastic
   c. Neither

Determining existing knowledge about collision types

Figure C.3: Momentum and collision lab pre-test

Momentum & Collision Lab Post-Test

Name :.. Question purpose in yellow (removed in the original version)

Group :..

Please work individually to answer the following questions:

1. Briefly, how would you define the following words, in terms of physics?

   Mass: ____

   Velocity: ____

   Momentum: ____

   Comparing definitions of key terms pre- and post-learning activity

2. Based on your work in the learning activity, which of the following is a valid formula for the relationship between mass (*m*), velocity (*v*) and momentum (*p*)?
   a. pm = v
   b. ½ p = mv
   c. $mv = p^e$
   d. p = mv

   Directly testing understanding of p = mv

2. A pool ball is rolling along when it collides with a stationary ball of the same weight. After the collision, both balls stick together and continue to travel at the same velocity, in the same direction as before. Compared to the original pool ball, has the group of two pool balls:
   a. Gained momentum
   b. No change in momentum
   c. Lost momentum
   d. It never had any momentum at al

   Testing ability to apply p = mv

2. The collision described in Question 1 is:
   a. Elastic
   b. Inelastic
   c. Explosive

   Testing knowledge about collision types

Figure C.4: Momentum and collision lab post-test

Pressure Lab Pre-Test

Name :... Question purpose in yellow (removed in the original version)

Group :...

Determining baseline definitions of fluids, compressibility & pressure

1. Briefly, what are your definitions of the following words, in terms of physics?

   Fluids: ____

   Compressible: ____

   Pressure: ____

Testing understanding of the relationship between volume and forces acting on compressible fluids

2. The main tank of an air compressor has a volume of 30L. When the compressor is turned on, a force is applied to the air to move it into a smaller holding tank. After this force is applied, will the air in the holding tank have a volume (circle one)

   greater than 30 L / less than 30 L / equal to 30L

Testing understanding of the relationship between pressure and depth

3. i) You're floating on an inner tube at the beach, when all of a sudden, a shark steals it from you and swims down deep into the ocean. What do you expect to happen to the inner tube several hundred feet below sea level?

   a) The tube's volume will increase and it will eventually pop
   b) The tube's volume will increase but it will not pop
   c) The tube's volume will not change
   d) The tube will crumple as its volume decreases

   ii) What is causing this change?

   ____

Testing understanding of the relationship between pressure and surface area

4. Jamal puts downward pressure of 200 Pa on his surfboard, which has an area of 4 m^2. However, the back of his surfboard is suddenly bit off by the same shark that took your inner tube! Now that the surfboard's area is reduced, how does the pressure Jamal places on the surfboard compare to the original pressure of 200 Pa?

   a) Jamal now places more than 200 Pa on the surfboard
   b) Jamal still places 200 Pa on the surfboard
   c) Jamal now places less than 200 Pa on the surfboard

Figure C.5: Pressure lab pre-test

Pressure Lab Post-Test

Name :.. Question purpose in yellow (removed in the original version)

Group :..

Determining whether definitions have changed pre- vs post-activity

1. Briefly, what are your definitions of the following words, in terms of physics?

   Fluids: ____

   Compressible: ____

   Pressure: ____

Check your understanding of the concepts of fluids, pressure and hydraulics:

Testing understanding of the relationship between volume and forces acting on compressible fluids

2. You have two bottles that are completely full, one with a gas and one with liquid. You place them both in a pressurized chamber that increases the air pressure. At the new pressure, do you think the volume of each bottle will be greater than, equal to, or less than the volume at normal air pressure?

   Water bottle: greater volume / equal volume / lesser volume (circle one)

   Helium bottle: greater volume / equal volume / lesser volume (circle one)

Testing understanding of the relationship between pressure and depth

3. a) You go on a long hike up a mountain for your friend's birthday. When you get to the top, you notice that the birthday balloon that you half-filled at the bottom is now fully inflated. Is the pressure from outside the balloon greater at the top of the mountain or the base? (Circle one)

   Pressure is greater at the top / Pressure is greater at the base / Pressure is equal

   b) Briefly explain your choice

   ____

Testing understanding of the relationship between pressure and surface area

4. Ai Ling goes ice skating. Her downward force from her body weight creates a pressure of 150 Pascals. The blades of her ice skates have a combined surface area of 50 cm^2. Later, she takes her ice skates off and walks around in her shoes (with combined surface area 200 cm^2). How does the pressure she exerts on the ground compare between shoes and skates?

   a) She exerts more pressure on the ground wearing skates
   b) She exerts equal pressure on the ground wearing skates vs. wearing shoes
   c) She exerts more pressure on the ground wearing shoes

Figure C.6: Pressure lab post-test