

DISS. ETH NO. 30445

USER AGENCY AND SYSTEM AUTOMATION
IN INTERACTIVE INTELLIGENT SYSTEMS.

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES
(Dr. sc. ETH Zurich)

presented by

THOMAS LANGERAK
M.Sc., Aalto University
M.Sc., University of Twente

born on 13 January 1994

accepted on the recommendation of

Prof. Dr. C. Holz
Prof. Dr. A. Oulasvirta
Prof. Dr. R. Murray-Smith

2024

Thomas Langerak: *User Agency and System Automation in Interactive Intelligent Systems*. © 2024

DOI: 10.3929/ethz-b-000705078

ABSTRACT

The balance between user agency and system automation in interactive intelligent systems is crucial for intuitive and efficient interactions. While fully automated systems could potentially offer greater efficiency and demonstrably improved performance, making them perfect is notoriously hard. The inevitable shortcomings of automated systems diminish usability and overall experience, thereby compromising users' perceived self-determination. Conversely, tools, systems that rely entirely on user agency and have no level of automation, though offering full control to the user, can be inefficient and fail to enhance the user's capabilities. Hence, for effective human-AI interactions, we need to find a balance between user agency and system automation. The question we address in this dissertation is "How can we balance user agency and system automation for the interaction with intelligent systems?"

We approach this challenge through four main contributions. First, we introduce a novel spherical electromagnet capable of generating adjustable forces on an untethered tool, allowing users to feel grounded forces while maintaining full agency. Second, we develop an integrated sensing and actuation system that tracks a passive magnetic tool in 3D space while simultaneously delivering haptic feedback, eliminating the need for external tracking. Third, we propose an optimal control method for electromagnetic haptic guidance systems that balances user input and system control, allowing users to adjust trajectories and speed as needed. Finally, we present a model-free reinforcement learning approach for adaptive user interfaces that learns interface adaptations without relying on heuristics or real user data.

Our findings, based on simulations and user studies, suggest that the shared control of intelligent systems has the potential to significantly outperform naive control strategies. Thus, we contribute methodologies that find an agency-automation trade-off and pave the way for more interaction with intelligent systems. Our research demonstrates that integrating models of human behavior, either explicitly or implicitly, into control strategies enables intelligent systems to better account for user agency. We show that the trade-off between user agency and system automation is not solely an algorithmic problem but must also be considered in the engineering of physical devices and interface design. We advocate for an integrated end-to-end approach to interaction with intelligent systems that incorporates algorithmic, engineering, and design perspectives.

ZUSAMMENFASSUNG

Das Gleichgewicht zwischen der Handlungsfreiheit des Benutzers und der Systemautomatisierung in interaktiven intelligenten Systemen ist entscheidend für intuitive und effiziente Interaktionen. Während vollautomatisierte Systeme potenziell größere Effizienz und nachweislich verbesserte Leistung bieten könnten, ist es bekanntermaßen schwierig, sie perfekt zu machen. Die unvermeidbaren Mängel automatisierter Systeme verringern die Benutzerfreundlichkeit und das Gesamterlebnis und beeinträchtigen dadurch die wahrgenommene Selbstbestimmung der Benutzer. Umgekehrt können Werkzeuge, d.h. Systeme, die sich vollständig auf die Benutzerautonomie stützen und kein Automatisierungsniveau aufweisen, zwar die volle Kontrolle für den Benutzer bieten, aber ineffizient sein und die Fähigkeiten des Benutzers nicht erweitern. Daher müssen wir für effektive Mensch-KI-Interaktionen ein Gleichgewicht zwischen Benutzerautonomie und Systemautomatisierung finden. Die Frage, die wir in dieser Dissertation behandeln, lautet: "Wie können wir Benutzerautonomie und Systemautomatisierung für die Interaktion mit intelligenten Systemen ausbalancieren?"

Wir nähern uns dieser Herausforderung durch vier Hauptbeiträge. Erstens stellen wir einen neuartigen sphärischen Elektromagneten vor, der in der Lage ist, einstellbare Kräfte auf ein kabelloses Werkzeug auszuüben und es den Benutzern ermöglicht, geerdete Kräfte zu spüren, während sie volle Autonomie behalten. Zweitens entwickeln wir ein integriertes Erfassungs- und Betätigungssystem, das ein passives magnetisches Werkzeug im 3D-Raum verfolgt und gleichzeitig haptisches Feedback liefert, wodurch die Notwendigkeit einer externen Verfolgung entfällt. Drittens schlagen wir eine optimale Kontrollmethode für elektromagnetische haptische Führungssysteme vor, die Benutzereingaben und Systemsteuerung ausbalanciert und es den Benutzern ermöglicht, Trajektorien und Geschwindigkeit nach Bedarf anzupassen. Schließlich präsentieren wir einen modellfreien Ansatz des verstärkenden Lernens für adaptive Benutzeroberflächen, der Schnittstellenanpassungen erlernt, ohne sich auf Heuristiken oder reale Benutzerdaten zu stützen.

Unsere Ergebnisse, basierend auf Simulationen und Benutzerstudien, deuten darauf hin, dass die geteilte Kontrolle intelligenter Systeme das Potenzial hat, naive Kontrollstrategien deutlich zu übertreffen. Somit tragen wir Methoden bei, die einen Kompromiss zwischen Autonomie und Automatisierung finden und den Weg für mehr Interaktion mit intelligenten Systemen ebnen. Unsere Forschung zeigt, dass die Integration von Modellen menschlichen Verhaltens, sei es explizit oder implizit, in Kontrollstrategien es intelligenten Systemen ermöglicht, die Benutzerautonomie

besser zu berücksichtigen. Wir zeigen, dass der Kompromiss zwischen Benutzerautonomie und Systemautomatisierung nicht nur ein algorithmisches Problem ist, sondern auch bei der Entwicklung physischer Geräte und dem Interface-Design berücksichtigt werden muss. Wir befürworten einen integrierten End-to-End-Ansatz für die Interaktion mit intelligenten Systemen, der algorithmische, technische und Design-Perspektiven einbezieht.

ACKNOWLEDGEMENTS

First and foremost, I would like to express my gratitude to Prof. Otmar Hilliges. He believed in me, despite my unconventional background, and gave me the opportunity to explore this academic world. Under Otmar's supervision and mentorship, I have grown into the researcher I am today. My PhD journey was a long learning experience, filled with both ups and downs. I am immensely grateful for his patience and trust throughout this process. The weekly one-on-one meetings with him were invaluable to my work, and I only came to appreciate them too late.

I would also like to extend my sincere thanks to my committee. Prof. Christian Holz, who was not only a close collaborator but also stepped in as my thesis supervisor when it became necessary, has my sincere gratitude. I am also deeply thankful to Prof. Antti Oulasvirta, who introduced me to the field of computational interaction and whose opinions I value. Antti ignited my initial passion for HCI and research. Finally, I would like to thank Prof. Roderick Murray-Smith, whose work was instrumental in shaping my understanding of interaction, and who graciously agreed to join my committee despite our limited prior interaction.

Furthermore, I want to thank all the external collaborators who played a part in this journey, including Prof. Daniele Panozzo, Dr. Bernhard Thomaszewski, Dr. Tanya Jonker, and Dr. Kashyap Todi. Their fresh perspectives were always invaluable.

A very special thanks goes to all my colleagues and friends at the AIT Lab. Together, we truly achieved more than we could have as individuals. You ensured that the lab was a place of great times, fun, and an enjoyable atmosphere. I had countless discussions with each of you that shaped my research. I want to specifically thank Dr. Juan Zárate for always engaging in conversations about science, magnets, and life. I am also grateful to Sammy Christen for patiently listening to my frustrations and for teaching me about Reinforcement Learning. My thanks also go to Dr. Velko Vechev and Dr. Christoph Gebhardt for their unwavering support, whether it was meeting deadlines, answering HCI-related questions, or help TAing. Marcel Buehler deserves thanks for making the office a great place to be, encouraging me to go bouldering, and for bringing me El Tony or a Früchtbrotli. I also want to thank Prof. David Lindlbauer for our beer-fueled discussions, his selfless can-do attitude during deadlines, and for enabling me to join the UIST Organization Committee. Finally, I am grateful to Prof. Xucong Zhang, Dr. Emre Aksan, Dr. Manuel Kaufmann, Yufeng Zheng, Prof. Jie Song, Prof. Anna Feit, Alex Fan, Mert Albaba, and all the others in the lab for being part of this journey.

I also want to express my gratitude to my friends for patiently listening to my frequent complaints, helping me switch off from work, and ensuring I stayed grounded in reality. A special thanks goes to my friends back in the Netherlands; I know it's not always easy to stay in touch with someone who has moved away, and I appreciate the effort they made. I also want to thank Salma Thalji for always being there, for providing me with pragmatic advice, for teaching me that motivation isn't a requirement to accomplish something, and for her understanding. Meeting you on that New York rooftop during the summer school is a memory I will always cherish.

Last but certainly not least, I want to thank my parents for supporting me in countless ways throughout my studies and for offering (sometimes unsolicited) advice when I needed it, even when I didn't realize I did.

CONTENTS

I Introduction

1	Introduction	1
1.1	Motivation	1
1.2	Approach	3
2	Contributions	7
2.1	Structure of Dissertation	9
2.2	Publications	10
3	Related Work	13
3.1	Human-AI Interaction	13
3.2	Control Theory in Interaction	14
3.3	Haptic Interfaces	17
3.4	Adaptive Interfaces	23
4	Background	27
4.1	Magnetic Model	27
4.2	Optimal Control	30

II The Design of Shared Variable Interfaces

5	Contact-free non-planar haptics with a spherical electromagnet	41
5.1	Introduction	42
5.2	System Description	44
5.3	System Evaluation	51
5.4	User Evaluation	53
5.5	Discussion	55
6	Volumetric sensing and actuation of passive magnetic tools for dynamic haptic feedback	57
6.1	Introduction	58
6.2	System Overview	59

6.3	Method	61
6.4	Evaluation	66
6.5	Applications	71
6.6	Discussion	73
7	Summary & Insights	75
III The Control of Shared Variables Interfaces		
8	Optimal control for electromagnetic haptic guidance systems	79
8.1	Introduction	80
8.2	Method Overview	83
8.3	Method	83
8.4	Implementation	89
8.5	Evaluation	94
8.6	Applications	98
8.7	Discussion	101
9	Multi-Agent Reinforcement Learning for Goal-Agnostic Adaptive User Interfaces	103
9.1	Introduction	104
9.2	Framework Overview	107
9.3	Method	109
9.4	Implementation	112
9.5	Evaluation	113
9.6	Applications	120
9.7	Discussion	125
10	Summary & Insights	127
IV Conclusion		
11	Conclusion	133
12	Outlook and Future Work	137
12.1	Understanding Humans	137
12.2	Understanding what Humans Understand	138
12.3	Cooperative Control	139

Bibliography 141

V Appendix

- A Optimal Control for Electromagnetic Haptic Guidance Systems 175
 - A.1 Dipole-Dipole Model 175

- B MARLUI: Multi-Agent Reinforcement Learning for Adaptive UIs 181
 - B.1 Learned Lower Level 181

Part I

INTRODUCTION

INTRODUCTION

1.1 MOTIVATION

Although human-computer interaction may appear relatively unchanged since the post-war era, as we still mainly rely on the mouse and keyboard, today's intelligent systems depend on much more than just explicit user input. Autonomous vehicles utilize computer vision to anticipate human intent and movement [96, 114, 174]. GitHub Copilot automatically adapts and generates code using large language models [107, 273]. Smart home devices, like Nest thermostats, learn user preferences to automatically optimize comfort and energy efficiency [212, 279], and advanced software like photo editing tools amplify human creativity [162]. These modern systems leverage contextual understanding to adapt to user needs. However, the shift from explicit to implicit user input, which comes from contextual understanding, is not without challenges. Systems that take action without explicit inputs reduce user agency. Thus, despite significant advances in the contextual understanding that these intelligent systems leverage, interaction with these systems remains challenging. In this dissertation, we focus on algorithmic methods to balance the automation provided by implicit contextual understanding with user agency achieved through explicit input in the area of Human-Computer Interaction (HCI).

In intelligent systems, a user interacts with an artificial agent that has contextual knowledge. The incorporation of contextual understanding aims to make these systems more natural and beneficial for real-world human tasks [269]. The goal of the artificial agent is to assist users in completing tasks, perform tasks on the users' behalf, or otherwise help them. In contrast to non-context-aware devices (e.g., [65, 66]) that rely on explicit and precise user input, intelligent systems are always available. This crucial property allows intelligent systems to proactively engage with their environment and users. In turn, this property enables implicit and imprecise input from the user. Humans leverage these kinds of interactions on a daily basis – e.g., saying "can you hand me *this*?" [147]. Furthermore, the proactive capabilities enable intelligent systems to act on behalf of the user by inferring their intent – e.g., a robot handing you a glass of water without being told to do so [49]. Yet, intelligent systems are still capable of acting based on explicit user commands, seemingly losing no ground to classical forms of human-machine interaction. However, as intelligent systems are capable of both implicit (system automation) and explicit

(user agency) interactions, they need to strike a balance between user agency and the automation of tasks.

This agency-automation trade-off becomes even more crucial in interactions where the system and user operate on a shared variable, that is, a variable that can be manipulated simultaneously by both the user and the system. For instance, consider the control of a semi-autonomous car. The user changes the car's acceleration, while the intelligent system simultaneously adjusts it. The artificial agent might increase safety and driving efficiency. However, it might not include passenger preferences. Similarly, a drawing assistance system might take control of a pen, thereby increasing accuracy yet constraining the user's creative expression. Also consider a virtual reality intelligent adaptive user interface. The interface may show optimal adaptations and decrease task completion time, but it does not necessarily align with user expectations. In all these scenarios, the ownership of the variable is ambiguous, and therefore, the control is ambiguous.

This ambiguity and agency-automation balance is where the primary challenge for interaction with intelligent systems lies. Fully automated systems could potentially provide greater efficiency and measurably better performance but at the cost of compromising user agency and diminishing perceived usability and overall experience. On the other hand, systems that rely solely on user agency provide complete control but may be sub-optimal, inefficient, and fail to augment the user's abilities. This dissertation explores the question: *How can we algorithmically control intelligent systems with shared variables to balance user agency and system automation?*

Researchers have explored the balance between agency and automation in intelligent systems. In the 1960s, Bar-Hillel [18] addressed optimizing human-computer labor division in translation tasks. The 1990s saw debates over user agency versus system automation [245], leading to a consensus on automation that enhances productivity while maintaining control. Additionally, research on human-AI collaboration emphasizes the importance of bidirectional communication, trust, and shared goals to enhance team effectiveness [61, 244]. To facilitate bidirectional communication, both models of human behavior and intelligent control strategies are necessary. Traditional methods to model human-computer interaction have focused on manual control theory [52, 180] using heuristics [35–37, 116, 127] and low-parameter mathematical models [74, 98]. More recent approaches leverage reinforcement learning (RL) due to its ability to predict user behavior [83, 118, 119]. These models of human behavior need to be integrated into the control strategies of systems. However, most haptic devices use a form of open-loop control [278], and thus do not take the user into account. Alternatively, some devices use proportional-integral-derivative control (PID) [3, 214, 219]. However, PID, among other shortcomings, optimizes for a fixed setpoint, thereby limiting user freedom. This dissertation builds upon

these foundations by integrating simple heuristics and RL to model human behavior in both the control of pen-based haptic feedback systems and the policy of adaptive user interfaces. We introduce novel kinesthetic haptic devices, essential for exploring the agency-automation spectrum, and employ model-based control strategies and multi-agent RL to dynamically adjust to user interactions, providing real-time optimal adaptations and enhancing user experience.

In summary, AI-driven systems are becoming increasingly ubiquitous in people's work and lives. These systems are no longer just "conventional" computing systems, but intelligent systems that exhibit unique characteristics that bring new interaction paradigms. This shift presents a challenge in balancing user agency and system automation, particularly when the user and system share control over a variable. This dissertation aims to explore how to strike an optimal balance between user control and system efficiency in interactions with intelligent systems and shared variables. We approach this question by introducing novel haptic devices, model-based control, and reinforcement learning approaches.

1.2 APPROACH

This dissertation aims to bridge the gap between user agency and system automation through four projects divided into two primary directions. First, we discuss the design of shared variable interfaces. Second, we explore the control strategies for such interfaces.

1.2.1 *The Design of Shared Variable Interfaces*

In shared variable interfaces, both the user and the system act on a shared variable, which can take many forms, such as the acceleration of a car or a graphical menu in an operating system. A special instance of this shared variable is kinesthetic haptic feedback. Here, not only is the variable itself shared (e.g., the position of a joystick), but the action and perception of the variable are in the same modality from the user's perspective. For example, to change or perceive a joystick's position, the user uses their hands. This singular interaction modality contrasts sharply with, for example, touch interfaces, where users perceive the UI through their eyes but interact with it using their hands. Using an interface with a single interaction modality potentially reduces confounding factors when investigating interactions with shared variable interfaces. This shared domain allows for simultaneous interaction between a user and an intelligent system.

Previous work in kinesthetic haptic feedback generally revolves around complex mechanical contraptions (e.g. [9, 177, 249, 253, 267, 291]). Further extensions of

such systems include exoskeletons [44, 92], gloves [55, 99], and tilt-platforms [128, 217]. However, these often require user instrumentation and inherently involve system friction, which users will always perceive. This friction leads to a system that, even when off, does not allow for full user agency.

To overcome the limitations of not covering the full spectrum from user agency to system automation and to enable our research, we introduce novel haptic interfaces. These haptic devices sense and actuate a passive tool. We introduce a novel spherical electromagnet that acts on an embedded permanent magnet, enabling an untethered tool that allows for full user agency without physical constraints while simultaneously providing large grounded forces, which enables system automation.

In Part II, we introduce i) a novel spherical electromagnet (Chapter 5, published in [286]), and ii) a gradient-based tool tracking algorithm (Chapter 6, published in [145]). This combination allows us to deliver dynamic haptic feedback. We evaluate our system both technically and with users. We find that interacting with a shared variable in the same modality allows for natural interactions that increase user agency. However, our findings also suggest that more intelligent system control is necessary to enable more complex dynamics.

1.2.2 *The Control of Shared Variable Interfaces*

Previous systems typically use open-loop control [278] to control intelligent systems. This approach fails to consider the user, eliminating the capability to trade off user agency with system automation. Alternatively, some systems use heuristics [33, 165, 250, 254], which require tediously crafted rules by experts. Furthermore, supervised learning [27, 69, 146, 173, 178, 181, 208, 238, 239] and multi-armed bandits [88, 122, 131, 134, 135, 163] are popular approaches for controlling intelligent systems. However, these approaches optimize for a myopic decision, failing to consider future states of both the system and the user. Optimizing purely for the next timestep limits the system’s capabilities to intelligently trade off user agency and system automation. In contrast, our work minimizes a cost function over a receding horizon[,] taking into account expected user and system behavior. This enables an explicitly optimized balance between agency and automation, taking into account possible future states.

In Part III, we first discuss Model Predictive Contour Control (MPCC), an optimization-based control strategy that uses a supplied model of the system dynamics in combination with a cost function (Chapter 8, published in [143]). We apply this in the context of a prototypical pen-based electromagnetic haptic-feedback system. Our approach allows users to easily override the system, adapt their input spontaneously, and draw at their own speed. While beneficial, crafting a system dynamic that includes user behavior is challenging. Secondly, we focus on UI adaptation as

a use case of an interface with a shared variable (Chapter 9, published in [142]). Moving from haptics to a digital-only interface offers several advantages. Adaptive User Interfaces allow us to focus on control strategies, overcoming any sim-to-real gap confounding factors. Furthermore, existing cognitive models (such as Fitts' Law [74]) can be utilized in this context. To overcome the challenge of crafting system dynamics that include user behavior, we turn to Reinforcement Learning (RL) as a control strategy and a means to learn both user and system dynamics models that are intertwined. By formulating UI adaptation as a multi-agent reinforcement learning problem, we introduce a user agent that mimics a real user and learns to interact with an interface, while simultaneously introducing an agent that learns a control strategy to maximize the user agent's performance. The control agent learns the task structure from the user agent's behavior and, based on that knowledge, can support the user agent in completing its task.

CONTRIBUTIONS

We summarize the four main contributions of our dissertation.

1. CONTACT-FREE NON-PLANAR HAPTICS WITH A SPHERICAL ELECTROMAGNET

Many kinesthetic haptic feedback devices rely on mechanical contraptions to provide forces. However, even when unactuated, these devices limit user agency, as the mechanical components restrict movement. To investigate the agency-automation trade-off, we need to design an *untethered* haptic feedback device capable of delivering grounded forces. To that end, **we introduce *Omni-v1* (Chapter 5), a novel contact-free volumetric kinesthetic haptic feedback device.** This system combines a spherical electromagnet with a dipole magnet model and a simple control law to deliver dynamically adjustable forces onto a handheld tool. We conducted a user experiment with 6 participants to characterize the force delivery aspects and perceived precision of our system.

2. VOLUMETRIC SENSING AND ACTUATION OF PASSIVE MAGNETIC TOOLS FOR DYNAMIC HAPTIC FEEDBACK

Many haptic devices (including *Omni-v1*) require external optical tracking, which is often expensive, cumbersome, and requires a clear line of sight. To overcome these limitations, we developed *Omni-v2* (Chapter 6). *Omni-v2* is a haptic feedback device with integrated spatial tracking. The spatial tracking capabilities of *Omni-v2* are enabled by **a novel gradient-based method, which reconstructs the 3D position of the permanent magnet in mid-air using measurements from eight off-the-shelf Hall sensors integrated into the base.** Furthermore, following improvements to the actuator design, *Omni-v2* delivers over twice the forces compared to *Omni-v1*. We detail *Omni-v2*'s hardware implementation and our 3D reconstruction algorithm, providing an in-depth evaluation of its tracking performance. *Omni-v2* shows how integrating sensing and actuation provides natural interaction that encourages user agency. Both *Omni-v1* and *Omni-v2* open up new applications in AR/VR, particularly in design, gaming, and object exploration.

3. OPTIMAL CONTROL FOR ELECTROMAGNETIC HAPTIC GUIDANCE SYSTEMS

Existing haptic devices typically employ open-loop control, which does not account for user feedback. Alternatively, they use proportional–integral–derivative (PID) control or heuristics, which are usually based on timed references, limiting user agency. **We propose *MagPen* (Chapter 8), a time-independent closed-loop control strategy that allows users to retain agency while receiving haptic guidance.** Our real-time approach assists in pen-based tasks such as drawing, sketching, or designing. By iteratively predicting the motion of an input device, such as a pen, and adjusting the position and strength of an underlying dynamic electromagnetic actuator, our method provides flexible guidance without diminishing user control. Experimental results demonstrate that our approach is more accurate and preferred by users compared to open-loop and time-dependent closed-loop methods.

4. MULTI-AGENT REINFORCEMENT LEARNING FOR POINT-AND-CLICK ADAPTIVE USER INTERFACES

Most control strategies rely on known system dynamics. However, in many HCI scenarios, users are integral to the system, and their behavior is complex and not easily modeled. To address this, we treat HCI as a multi-agent problem where an agent learns system, task, and user dynamics through interaction with a synthetic user agent. Specifically, **we introduce *MARLUI* (Chapter 9), a multi-agent reinforcement learning approach for adaptive user interfaces (AUIs).** In our formulation, a user agent mimics real user interactions with the UI, while an interface agent learns to adapt the UI to maximize the user agent’s performance. Our method captures the underlying task structure, system dynamics, and user behavior. Experiments show that the learned policies generalize well to real users and achieve performance comparable to data-driven supervised learning baselines.

2.1 STRUCTURE OF DISSERTATION

We describe the structure of this doctoral dissertation in the following. After establishing the state of the art and background knowledge on the relevant methods, we outline our main contributions. More specifically:

I. INTRODUCTION

- CHAPTER 3 introduces the related work on haptic devices, adaptive interfaces, user modelling and control strategies.
- CHAPTER 4 provides theoretical and mathematical background in electromagnetism and control theory.

II. THE DESIGN OF SHARED VARIABLE INTERFACES

- CHAPTER 5 **Contribution 1:** *Contact-free Non-Planar Haptics with a Spherical Electromagnet.*
- CHAPTER 6 **Contribution 2:** *Volumetric Sensing and Actuation of Passive Magnetic Tools for Dynamic Haptic Feedback.*
- CHAPTER 7 summarizes, discusses the implications, and highlights the limitations of our work on shared variables interfaces.

III. THE CONTROL OF SHARED VARIABLE INTERFACES

- CHAPTER 8 **Contribution 3:** *Multi-Agent Reinforcement Learning for Point-and-Click Adaptive User Interfaces.*
- CHAPTER 9 **Contribution 4:** *Optimal Control for Electromagnetic Haptic Guidance Systems.*
- CHAPTER 10 summarizes, discusses the implications, and highlights the limitations of our work on joint control of shared variables.

IV. CONCLUSION

- CHAPTER 11 summarizes our contributions.
- CHAPTER 12 highlights possible future research directions. and control theory.

2.2 PUBLICATIONS

The contributions of this thesis are based on the following publications, in order of appearance:

1. Zarate, J.*, **Langerak, T.***, Thomaszewski, B. & Hilliges, O. *Contact-free nonplanar haptics with a spherical electromagnet* in *2020 IEEE Haptics Symposium (HAPTICS)* (2020), 698 ¹
2. **Langerak, T.***, Zarate, J.*, Lindlbauer, D., Holz, C. & Hilliges, O. *Omni: Volumetric sensing and actuation of passive magnetic tools for dynamic haptic feedback* in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (2020), 594 ²
3. **Langerak, T.**, Zárate, J., Vechev, V., Lindlbauer, D., Panozzo, D. & Hilliges, O. *Optimal control for electromagnetic haptic guidance systems* in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (2020), 951 ³
4. **Langerak, T.**, Christen, S., Albaba, M., Gebhardt, C., Holz, C. & Hilliges, O. *MARLUI: Multi-Agent Reinforcement Learning for Adaptive Point-and-Click UIs* in *Proc. ACM Hum.- Comput. Interact.* 8, *EICS, Article 253* (2024) ⁴

Further publications that were conducted during the course of my PhD research but are out of scope of this thesis are listed below, in order of publication:

1. **Langerak, T.**, Zarate, J., Vechev, V., Panozzo, D. & Hilliges, O. *A demonstration on dynamic drawing guidance via electromagnetic haptic feedback* in *Adjunct Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (2019), 110 ⁵
2. Abler, A., Zárate, J., **Langerak, T.**, Vechev, V. & Hilliges, O. *Hedgehog: Handheld spherical pin array based on a central electromagnetic actuator* in *2021 IEEE World Haptics Conference (WHC)* (2021), 133 ⁶

1 My specific contribution in this work is the ideation, fabrication, control strategy, technical evaluation and user study.
 2 My contribution in this work is the ideation, fabrication, method and evaluation.
 3 I have been fully responsible for this work and contributed to all components, with exception of the derivation of the electromagnetic model we use.
 4 I have been fully responsible for this work and contributed to all components.
 5 I have been fully responsible for this work and contributed to all components.
 6 I have contributed with ideation, discussions and to the writing of the paper. The first author was a Master Thesis student under my supervision

3. **Langerak, T.**, Christen, S., Feit, A. M. & Hilliges, O. Generalizing User Models through Hybrid Hierarchical Control. *CHI Workshop on Reinforcement learning for Human-Computer Interaction* (2021) ⁷
4. Wu, M.*, **Langerak, T.***, Hilliges, O. & Zarate, J. Robust Real-Time Tracking of Axis-Symmetric Magnets via Neural Networks. *Transactions on Magnetics* (2024) ⁸
5. **Langerak, T.**, Todi, K., Lafreniere, B., Desai, R. & Jonker, T. XAIUI: Mental-Model Aware Explainable AI for Intelligent User Interfaces. *Under Review for Transactions on Interactive Intelligent Systems* (2024) ⁹
6. Albaba, M., Christen, S., Gehbhardt, C., **Langerak, T.**, Black, M. & Hilliges, O. RILe: Reinforced Imitation Learning. *Under Review for NeurIPS* (2024) ¹⁰

⁷ I have been fully responsible for this work and contributed to all components.

⁸ I have contributed with ideation, discussions and to the writing of the paper. The first author was a Master Thesis student under my supervision

⁹ I have been fully responsible for this work and contributed to all components.

¹⁰ I have contributed with fruitful discussions and to the writing of the paper.

RELATED WORK

We provide an overview of related work. We start by discussing human-AI interaction from a high-level perspective. Then, we will discuss control theory in the context of HCI, including, a brief overview of computational user modeling and its role in the control loop. Finally, we focus on two types of interfaces with shared variables, discussing what they are and specific examples of control. First, we discuss haptic interfaces, and secondly, we discuss adaptive user interfaces.

3.1 HUMAN-AI INTERACTION

Research on human-AI collaboration has revealed that integrated teams of humans and AI systems are more effective than either working independently [16, 17, 60]. Research explores various aspects of this collaboration, such as conceptual architectures and frameworks [117, 172, 196, 216], and performance metrics [16].

To gain insights into human-machine collaboration, researchers have applied psychological theories of human teamwork [68, 186]. These theories highlight core principles for bidirectional communication, trust, goals, situational awareness, language, intentions, and decision-making between humans and AI systems [61, 244]. This approach marks a shift from the traditional one-way communication model in conventional HCI contexts [269], emphasizing a reciprocal relationship. In this dissertation, we embed this reciprocal relationship by modeling human-AI interaction as a multi-agent problem. Furthermore, having a shared variable inherently forms a bidirectional communication channel.

There are foundational concepts in interpersonal teaming that require adaptation for application to human-machine teams. These include Autonomy (compared to Automation), Trust (compared to Reliability), and Teaming (compared to the Use of Automation) [91]. This adaptation has become particularly important in the context of autonomous car research [13, 151] and military applications [40]. Sycara & Lewis [259] identified three general roles that machines can support within teams: assisting individuals with their tasks, acting as equal team members, or supporting the team as a whole. This dissertation focuses on the first role, supporting individuals. Usually, this role is investigated in the context of decision support systems [168]. However, in contrast to decision support systems, our focus is on optimizing the control of interfaces to best support user interactions.




FIGURE 3.1: Human-Computer Interaction as a closed-loop system, adapted from [190]

3.2 CONTROL THEORY IN INTERACTION

Control theory deals with the behavior of dynamic systems with inputs and how their behavior is modified by feedback [20, 79]. The goal of control theory is to develop methods for influencing the behavior of dynamical systems to achieve desired outcomes [79]. The key idea is to measure the system’s output, compare it to a reference or desired state, and adjust the system’s input accordingly to minimize the difference between the actual and desired outputs [235]. This feedback loop allows the system to adapt and maintain stability in the presence of disturbances or uncertainties [79, 169]. Control theory provides a mathematical framework for modeling, analyzing, and optimizing such feedback systems [79, 246]. In an HCI context, a human interacts with a machine. While the human interaction can be modeled as a control system, the machine uses control theory to achieve desired outcomes given the human input. Together they form a closed-loop system (Figure 3.1) [190]. First, we will discuss control theory from a machine perspective. Then, we will discuss modeling human behavior from a control theory perspective. For a mathematical introduction, see the background chapter (Section 4.2).

3.2.1 *Control of Machines*

There has been extensive research on optimal control techniques and their applications to various domains. Proportional-integral-derivative (PID) control is one of the most widely used approaches due to its simplicity and effectiveness. For example, PID control has been used to control haptic devices [3, 214, 219]. However, PID control has some shortcomings, such as poor adjustment and follow-up when the

system has strong interference or high nonlinearity and uncertainty [136]. It can also struggle with precise control and overshoot in some applications [252].

Model Predictive Control (MPC) is another optimal control technique that has gained popularity in recent years. MPC addresses some of the limitations of PID control by explicitly taking into account the nature of the actuators and system dynamics over a certain time horizon before deciding the control action to be applied [57, 218]. This allows MPC to handle constraints, nonlinearities, and uncertainties more effectively. MPC has been used, for instance, in computer numerical controlled machines [140, 141], the control of drones [84], or autonomous vehicles [160]. However, MPC also has some shortcomings, such as the need for an accurate system model and the computational cost associated with solving the optimization problem in real-time [100]. Furthermore, MPC is a time-dependent control strategy, in which the setpoint moves at a fixed pace. This limits user creativity and freedom.

Reinforcement Learning (RL) is an increasingly popular approach for optimal control in intelligent systems that interact with humans. RL addresses some of the limitations of MPC by learning control policies directly from interaction with the environment, without requiring an explicit system model [257]. This allows RL to handle complex, uncertain, and nonlinear systems more effectively. RL has been applied to robotics, where it has been used for learning complex behaviors such as manipulation and locomotion [46, 48]. In self-driving cars, RL has been employed for learning driving policies that can handle diverse traffic scenarios and road conditions [130].

In this dissertation, we look at a time-independent MPC for haptic devices and RL for adaptive user interfaces. In both scenarios, we embedded user dynamics into the control loop. To that end, we must have models of human behavior.

3.2.2 *Human Behavior as Control Strategy*

The specific area of control systems relating to human users became a major focus starting in the 1950s [52, 180]. Our discussion of the relevant work is based on [190]; we refer to that for a more extensive discussion. In general, it is accepted that models of human behavior enable simulation. Simulating human behavior, in turn, has many advantages in theory crafting, design, novel interaction paradigms, and safety [191].

Manual control theory [52, 180], which seeks to model the interaction of humans with machines like aircraft or cars, grew out of Craik's early, war-related work [53]. It became more well-known in the broader framing of Wiener's Cybernetics [276]. According to Wickens and Hollands [275], modeling human control behavior emerged from two main schools: the skills researchers, who focused on learning and acquisition in undisturbed environments, and the dynamic systems

or manual control theory researchers [123, 241], who modeled the interaction of humans with machines in vehicle control and complex industrial process control, driven by engineering motivations to eliminate error in closed-loop systems. Poulton [215] reviews the early tracking literature, while Jagacinski and Flach [113] provide an accessible textbook review of manual control approaches.

Early work relies on heuristics [35–37, 116, 127] and on low-parameter mathematical models [74, 98]. More recent work extends these models and, for instance, predicts the operating time for a linear menu [50], gaze patterns [227], pointing [106, 176, 188, 192], or cognitive load [64].

Recently, reinforcement learning gained popularity within the research area of computational user models. This popularity is due to its neurological plausibility [30, 78], allowing it to serve as a model of human cognitive functioning. The underlying assumption of RL in HCI is that users behave rationally within their bounded resources [85, 203]. There is evidence that humans use such a strategy across domains, such as in causal reasoning [62] or perception [86]. In human-computer interaction, researchers have leveraged RL to automate the sequence of user actions in a keystroke-level-model framework [150] or to predict fatigue in volumetric movements [38]. It was also used to explain search behavior in user interfaces [280] or menus [42] and as a model for multitasking [119]. Most similar to our work is research on hierarchical reinforcement learning for user modeling. Jokinen *et al.* [118] show that human-like typing can emerge with the help of Fitts' Law and a gaze model. Other works show that HRL can elicit human-like behavior in task interleaving [83] or touch interactions [118].

Specifically, Multi-Agent Reinforcement Learning (MARL) is of interest [59]. MARL is interesting as it captures the implicit closed-loop iterative nature of HCI. We can see both the human and computer as two agents that interact together in an environment. MARL is popular in robotics [199]. Furthermore, MARL is closely related to alignment theory as MARL deals with training multiple AI agents to behave in desired ways. The challenge of aligning multiple agents' objectives with each other and with human values relates closely to fundamental alignment problems [170, 225]. MARL is also highly relevant to cooperative AI, as it provides frameworks for training multiple agents to work together toward common goals. Techniques from MARL, like centralized training with decentralized execution [39], can be used to develop AI systems that cooperate effectively [260].

In this dissertation, we first use simple heuristics to model human behavior in a pen-based haptic feedback guidance system. We embedded this human behavioral model into an MPC that controls the haptic system. Afterwards, we use RL to model human behavior in adaptive interfaces. This is used to implicitly learn a behavioral model by the adaptive interface, to show optimal adaptations to the user.

3.2.3 Shared Control

Traditionally, there is a clear distinction between assistance systems, where the machine only supports the human, and automation, where the machine takes over the main task, replacing the human. Sheridan recognized that this distinction should not be so black and white, and proposed levels of automation [242]. There are many situations where both the human and the machine should act together simultaneously, and where authority and tasks need to be shifted or adapted [182, 240]. Human-machine cooperation is a broader concept that includes shared control but also extends to cooperation at higher tactical and strategic levels like guidance and navigation. Shared control, specifically, is an approach where control is shared between human users and automated components, such as robots or computer systems. It involves cooperation at the control level between humans and machines [75]. Shared control is often discussed in the context of physical systems, such as (semi-)autonomous vehicles [175] and human-robot interaction [166].

3.3 HAPTIC INTERFACES

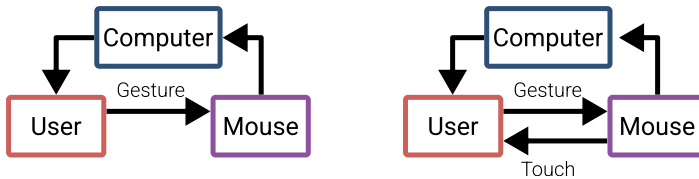


FIGURE 3.2: The Haptic-Loop allows for bi-directional exchange of information between machine and user (Hayward *et al.* [95]).

Human haptic perception, one of our primary senses, is spread throughout the body in the skin, muscles, and joints. It allows us to perceive our environment, both natural and simulated, through two distinct modes: tactile and kinesthetic feedback [54]. Tactile sensing provides information about pressure, shear forces, and vibrations via mechanoreceptors in the skin. Kinesthetic feedback, closely tied to proprioception, enables us to sense the position of our limbs, body posture, and large-scale forces acting on our hands and body through muscle and joint activations. Haptic perception is unique in that it actively involves mechanical interaction - we manipulate the environment through touch and force, using the resulting feedback to continuously adjust our actions.

A haptic interface facilitates a two-way flow of information, allowing the user and machine to exchange data simultaneously, setting it apart from haptic displays

and related graphical displays that only provide one-way information transfer. The haptic interface provides haptic feedback based on the user's input, creating a "haptic loop." For example, a standard mouse (left in Figure 3.2) can be enhanced to provide haptic feedback (right in Figure 3.2) that resists movement when the user approaches a disabled on-screen action. Hayward *et al.* [95] further describe this information exchange as an exchange of energy, a concept that can be applied when using physics-based optimization techniques to develop haptic interfaces.

Haptic interfaces are especially interesting as a use case in our settings of shared variables. First, haptic technologies have found renewed interest in the human-computer interaction community due to the recent reemergence of Augmented and Virtual Reality systems. Second, the physical nature of haptic interaction makes variable ownership more impactful.

3.3.1 *Devices*

3.3.1.1 *Mechanical haptic feedback*

The most common form of rendering haptic feedback has been vibrotactile actuators. They are popular in commercial devices, such as the controllers of game consoles, AR and VR controllers, and mobile phones. They can also be embedded directly into displays [272] or in clothing [56, 195].

Vibrotactile actuators usually produce coarse and global feedback, especially when used in handheld controllers (e. g. rendering touch contact in VR [19, 45]). Researchers have investigated how to overcome this limitation, such as by rendering interpolations between several such motors [110] or strategically distributing them across a controller (e. g. placing them under the fingers to render local grasp feedback [148]).

Alternatives to vibrotactile feedback often involve more complex articulated haptic elements, such as arms and braking mechanisms (e. g. [9, 177, 249, 253, 267, 291]). In general, these types of systems can provide local haptic feedback at higher levels of fidelity and can render both tactile and kinesthetic feedback. Further extensions of such systems are exoskeletons [44, 92], gloves [55, 99], and tilt-platforms [128, 217]. These platforms are usually (rigidly) anchored and can therefore supply large forces. These approaches, except DextrES [99], rely on mechanical structures and anchoring to the environment. Therefore, their use is mostly limited to high-end applications such as teleoperation.

3.3.1.2 *Contact-free haptic feedback*

A second line of research focuses on contact-free haptics, which provide rich and strong feedback, and overcome the need for expensive and complex mechanical setups [31].

Within the contact-free domain, many different actuation devices have been explored, e. g. ultra-sound pressure waves [102], active control of stylus motions [126], and drones [97]. The most popular and practical actuators in this domain use magnetism. The simplest form of magnetism is delivered by passive magnets that are embedded into interactive objects (e. g. [278]). The recent advance of consumer 3D printing has allowed this approach to actuate objects with arbitrary shape and function [197, 290]. A big shortcoming of passive magnets, however, is the lack of dynamic control over them and thus the forces users perceive during interaction.

3.3.1.3 *Electromagnetic haptic feedback*

The shortcoming of passive magnets can be addressed using electromagnetism and computational control of magnetic forces. Two-dimensional arrays of electromagnets can be combined with passive magnets that are worn [4, 21, 23, 24, 271, 278, 288] or embedded in tools and interactive objects [2, 120, 270]. Pangaro et al. [205] model the force-field of each electromagnet and combine these using standard aliasing techniques, allowing directed movement of multiple objects on the surface. Similarly, Yoshida et al. [282] use linear induction motors to control objects on a tabletop. Strasnick et al. [255] use electromagnets to control an object on a mobile phone case. Suzuki et al. [258] combine these two works and use a grid of electromagnetic cores to move objects on a tabletop. The actuation area can be increased by attaching an electromagnet to a biaxial linear stage [143, 144].

Similarly, by leveraging the electromagnetic forces in a coil between two permanent magnets, large and grounded forces can be delivered onto a joystick [22]. The main drawback of this approach is the requirement for a mechanical connection, limiting the range of motion and impeding contact-free haptics. Senkal et al. [232, 233] and Li et al. [152] propose to use magnetorheological fluids in joysticks. With the help of an electromagnetic field, the internal friction can be significantly increased. This allows for a large breaking force; however, it does not allow adding energy to the system.

Perhaps the most well-known haptic interface that builds on Lorentz forces is the Butterfly Haptics Maglev [25, 26]. The design of this system consists of a flotor bowl with six integrated coils to which an interaction handle is rigidly attached. This flotor bowl is levitated between magnets assemblies that are part of a stator bowl. Due to the Lorentz levitation, haptic feedback can be achieved in a degree of rotary

movement and also a small translation. Due to the small movements allowed, the device is mainly suitable for small-scale manipulations, e.g., where only the fingers are used.

The design of our system is fundamentally different. In contrast, *Omni-v1* and *Omni-v2* use a single omnidirectional electromagnet, thereby largely increasing the rotary capabilities. Furthermore, by omitting the levitating flotor bowl, our devices greatly reduce the complexity and therefore make it easier to fabricate, thereby more likely to foster adaptation. Finally, our designs allow for contact-free haptics, which is not possible in the Butterfly system.

Closely related work to *Omni-v1* and *Omni-v2* also includes Omnimagnet by Petruska et al. [211] and its variants [108]. Similar to our devices, their system generates an omnidirectional magnetic field. Their design differs by using three nested cuboid-shaped coils, causing force decay as the user moves along the surface of the device as well as an obstruction of heat dissipation. This limits the maximal strength and duration of actuation [67], and also makes the device only suitable for rendering vibrotactile stimuli [289]. Due to the cuboid shape, the center-to-center distance between the electromagnet and the permanent magnet is not constant among the surface, which causes high variance in the forces perceived. In contrast, *Omni*'s design is spherical, symmetrical, and has intertwined coils. This results in better heat dissipation and less variance in the force, thereby arguably improving the user experience.

3.3.1.4 *Haptic Guidance Devices*

Providing haptic guidance to users can provide benefits for learning [263] and short-term performance (cf. Abbink et al. [1]). Teranishi et al. [263] demonstrate that participants showed improved learning for handwriting skills when receiving guidance through a 3-DOF Phantom Omni device. Mullin et al. [189] use a similar device as a handwriting aid for rehabilitation. Forsyth and MacLean [77] show that force cues are beneficial in navigation tasks. The focus of these works is that users receive tight guidance (i.e., they are supposed to follow the system as closely as possible).

There exists a large range of devices and systems that aim at providing guidance to users. Comp*Pass [194] uses pantograph-like devices to assist users in drawing, while I-Draw [71] is a motorized drawing assistant. Lin et al. [157] use a magnet mounted on a small robotic arm to retain the correspondence between the pen and a portable base. Digital rubbing employs a comparable system using a solenoid for tracing over digital images on real paper [129]. While users handle larger-scale motions of the devices, they generally aim at having full control over the resulting drawing. Users can take back this control; however, these systems do not provide

a way to guide users back to the target trajectory. Besides the aforementioned systems, several works aid users in the process of crafting and manufacturing (cf. Zoran et al. [293]). Free-D [292] and D-Coil [209] assist users in sculpting physical artifacts by guiding them on a predefined 3D shape. Shilkrot et al. [243] propose an augmented paintbrush to assist users in painting. While users can override these systems to deviate from the target shape, they have no mechanism that guides users back to the target.

3.3.2 Control of Haptic Devices

To effectively control haptic devices, a closed loop needs to be formed. This means that the user needs to be sensed, and based on this inference, an actuation decision needs to be made. We first discuss magnetic sensing and then the control for haptic guidance devices.

3.3.2.1 Magnetic Sensing

Permanent magnets have been used for tracking objects in 3D, ranging from styli and other interactive objects [137, 155], jewelry [10] all the way to fingers [94], joints, and other biological tissues [28, 261].

Ample research exists on tracking permanent magnets. Most of the existing literature uses isotropic (i.e., spherically) shaped permanent magnets, because the dipole model most accurately resembles these [111]. However, some work also uses electromagnets attached to fingertips [43]. Due to this, the fingers can be tracked individually. One of the biggest challenges is that a closed-form solution of the magnet states (e.g., position) is unlikely to exist in most scenarios. Therefore, optimization-based methods are commonly used [228, 262] and more recently also neural networks [226]. However, these methods were often employed offline, suffered from large latency, or converged to local minima.

A key related work is Magnetips by McIntosh et al. [179]. They use a permanent magnet attached to a fingertip to interact with a watch. The permanent magnet is tracked around the watch and also used to provide vibrotactile haptic feedback. Magnetips multiplexes actuation and sensing. However, this causes significant delays (2 ms for every swap between tracking and actuating), which may pose an issue for scenarios that require continuous interactions.

The work that most resembles our work from an algorithmic point of view is [262]. They track multiple spherical magnets online using an analytical Jacobian, allowing solving with a quasi-Newton method. In contrast, *Omni-v2* tracks a single magnet while compensating for drastically changing electromagnetic fields, rather than tracking multiple permanent magnets in a static environment. We do this by adjusting

the dipole model so that it is suitable for electromagnetism. We also propose a novel formulation of the position reconstruction problem and an implementation in PyTorch that can leverage the framework’s auto-grad capabilities, thus avoiding the need for an analytical Jacobian.

3.3.2.2 *Haptic Guidance Control*

Closest to *MagPen* in terms of hardware is dePENd by Yamaoka et al. [278]. They move a permanent neodymium magnet on a two-axis setup to control the metal tip of a ballpoint pen. The neodymium magnet "drags" the input pen around a predefined path, similar to a plotter. dePENd employs an open-loop strategy to control the magnet, which means users cannot deviate from the predefined path without risking losing haptic guidance. In contrast, in *MagPen* we propose a mathematical model and optimal control strategy that allows users to move at their own pace through a drawing, for example, and reacts in real-time to user input by altering the position and strength of the magnet. We show that our approach provides better results than their open-loop approach, as well as existing closed-loop approaches.

Kianzad et al. [125] use a ballpoint drive to assist users in sketching. They employ a proportional-derivative (PD) control loop, which allows users to deviate from the target to a certain extent. We show in our experiments that our optimization scheme outperforms such existing closed-loop approaches. Muscle-Plotter [165] proposes active guidance for users based on electrical muscle stimulation. Their control strategy is based on heuristics for users to share control with the system. Our approach could be applied to their work if the electromagnetic force model is replaced by a model of the interaction between the muscle stimulation and the force users produce.

Optimal reference following given real-world influences is studied in depth in the control theory literature. Methods like Model Predictive Control (MPC) [70] optimize the reference path and the actuator inputs simultaneously based on the system state. MPC is widely applied in many robotics (e.g., to control quadcopters, Mueller et al. [187]) and graphics applications (e.g., for human motion prediction, Da Silva et al. [57]). However, Aguiar et al. [5] show that the tracking error for following timed trajectories can be larger than if following a geometric path only.

To address this issue, Lam et al. propose Model Predictive Contouring Control (MPCC) [141] to follow a time-free reference, optimizing system control inputs for time-optimal progress. MPCC has been successfully applied in industrial contouring [141], RC racing [159], and in drone cinematography [193].

We also pose our optimization problem in this well-established framework. However, to the best of our knowledge, we are the first to apply it for haptic guidance systems where one has to consider both a controllable (i.e., the electromagnetic

force) and non-controllable (i.e., the user) system. We contribute a formulation of the problem including models and control algorithms to enable employing MPCC in this context.

3.4 ADAPTIVE INTERFACES

UI adaptation can either be offline, to computationally design an interface, or online, to adapt the UI according to users' goals. We will focus on online adaptive UIs and refer readers to [201, 202] for an overview of computational UI design.

3.4.1 *Control of Adaptive Interfaces*

We introduce an overview of different control strategies for adaptive interfaces. Note that we use control strategy and policy in the context of adaptive interfaces interchangeably.

3.4.1.1 *Heuristics, Bayesian Networks & Combinatorial Optimization*

In early works, heuristic- or knowledge-based approaches are used to adapt the UI [33, 250, 254]. Similarly, multi-agent systems employ rule-based and message-passing approaches [221, 222, 281]. Another popular technique for AUIs is domain-expert-designed Bayesian networks [29, 101]. More recently, combinatorial optimization was used to adapt interfaces dynamically [158, 206]. The downside of these approaches is that experts need to specify user goals using complex rule-based systems or mathematical formulations. Creating them comprehensively and accurately requires developers to foresee all possible user states, which is tedious and requires expert knowledge. Commonly, these approaches also get into conflict when multiple rules or objectives apply. This conflict often results in unintuitive adaptations. In contrast, *MARLUI* only requires the layout of the UI. From its representation as an RL environment, we learn policies that meaningfully adapt the UI and realistically reproduce user behavior.

3.4.1.2 *Supervised Learning*

Leveraging machine learning can overcome the limitations of heuristic-, network-, and optimization-based systems by learning appropriate UI adaptations from user data. Traditional machine learning approaches commonly learn a mapping from user input to UI adaptation.

Algorithms like nearest neighbor [146, 173], Naïve Bayes [69, 178], perceptron [238, 239], support vector machines [27], or random forests [181, 208] are used

and models are learned offline [27] and online [238]. Due to the problem setting, these approaches require users' input to be highly predictive of the most appropriate adaptation. Furthermore, it restricts the methods to work in use cases where myopic planning is sufficient, i.e., a single UI adaptation leads users to their goal. In contrast, *MARLUI* considers multiple goals when selecting an adaptation and can lead users to their goal using sequences of adaptations.

More recent work overcomes the limitations stemming from simple input-to-adaptation mapping by following a two-step approach. They (1) infer users' intention based on observations and (2) choose an appropriate adaptation based on the inferred intent [204]. Such work uses neural networks, and user intention is modeled either explicitly [132, 251] or as a low-dimensional latent representation [223].

However, these approaches are still highly dependent on the quality of the training data, which may not even be available for emerging technologies. In contrast, *MARLUI* can learn supportive policies without pre-collected user data by just observing simulated user behavior.

3.4.1.3 *Bandits & Bayesian Optimization*

Bandit systems are a probabilistic approach often used in recommender systems [88]. In a multi-armed bandit setting, each adaptation is modeled as an arm with a probability distribution describing the expected reward. The Bayes theorem updates the expectation, given a new observation and prior data. Related work leverages this approach for AUIs [122, 131, 163].

Bayesian optimization is a sample-efficient global optimization method that finds optimal solutions in multi-dimensional spaces by probing a black box function [236]. In the case of AUIs, it is used to find optimal UI adaptations by sampling users' preferences [134, 135]. Both approaches trade off exploration and exploitation when searching for appropriate adaptations (i.e., exploration finds entirely new solutions, and exploitation improves existing solutions), rendering them suitable approaches to the AUI problem. However, such methods are not able to plan adaptations over a sequence of interaction steps, i.e., they plan myopic strategies.

In addition, these approaches need to sample user feedback to learn or optimize for meaningful adaptations and, hence, also rely on high-fidelity user data. Furthermore, as users themselves learn during training or optimization, solutions can converge to sub-optimal user behavior as such methods reduce exploration with convergence. In contrast, *MARLUI* can plan adaptations over a sequence of interaction steps learned from realistic, simulated user data.

3.4.1.4 *Reinforcement Learning*

Reinforcement learning is a natural approach to solving the AUI problem, as its underlying decision-making formalism implicitly captures the closed-loop iterative nature of HCI [104]. It is a generalization of bandits and learns policies for longer horizons, where current actions can influence future states. This generalization enables selecting UI adaptations according to user goals that require multiple interaction steps. Its capability makes RL a powerful approach for AUIs with applications in dialog systems [80, 256], crowdsourcing [58, 105], sequential recommendations [41, 156, 161], information filtering [234], personalized web page design [72], and mixed reality [82]. Similar to our work is a model-based RL method that optimizes menu adaptations [266].

Current RL methods sample predictive models [80, 105, 266] or logged user traces [82]. However, these predictive models and offline traces represent user interactions with non-adaptive interfaces. Introducing an adaptive interface will change user behavior; so-called co-adaptation [171]. Hence, it is unclear if the learned model can choose meaningful adaptations when user behavior changes significantly due to the model's introduction. In contrast, our user agent learns to interact with the adapted UI; hence, our interface agent learns on behavioral traces from the adapted setting.

3.4.1.5 *Multi-Agent Reinforcement Learning*

MARL is a generalization of RL in which multiple agents act, competitively or cooperatively, in a shared environment [287]. Multi-agent systems are common in games [14, 112], robotics [109, 200], or modeling of social dilemmas [149, 284]. MARL is challenging since multiple agents change their behavior as training progresses, making the learning problem non-stationary. Common techniques to address this issue are via implicit [265] or explicit [76] communication, centralized critic functions [167], or curricula [164, 268].

BACKGROUND

We introduce background in electromagnetic models, specifically the dipole and dipole-dipole model. Afterwards, we introduce optimal control theory.

4.1 MAGNETIC MODEL

4.1.1 *Dipole Model*

In our haptic applications, we employ electromagnetic feedback. To effectively provide this feedback, it is crucial to understand the relationships between the electrical current, the magnetic field, and the force experienced by a permanent magnet embedded in a tool.

We utilize simplified magnetic field models from previous research [264], which allow us to simulate the magnetic field in real-time. These models require only the magnitude, orientation, and position of each dipole magnet. We apply similar simplifications in our approach.

The formula for the magnetic field created by a dipole is:

$$\mathbf{B}(\mathbf{r}, \mathbf{m}) = \frac{\mu_0}{4\pi} \frac{3 \hat{\mathbf{r}} (\mathbf{m} \cdot \hat{\mathbf{r}}) - \mathbf{m}}{|\mathbf{r}|^3}, \quad (4.1)$$

where μ_0 represents the permeability of air, \mathbf{r} is the vector from the center of the magnetic source to the testing point, and \mathbf{m} symbolizes the magnetic moment containing the *strength* and orientation of the magnetic source. The magnetic moment varies depending on the magnet type: \mathbf{m}_e for an electromagnet and \mathbf{m}_p for a permanent magnet. For the permanent magnet embedded in the tool, angles θ and φ are measured from the positive z-axis, as depicted in Figure 4.1.

The three components of the magnetic moment for the tool's embedded magnet are described by:

$$\mathbf{m}_p = \frac{1}{\mu_0} * Br * V * \begin{bmatrix} \sin(\theta) \cos(\varphi) \\ \sin(\theta) \sin(\varphi) \\ \cos(\theta) \end{bmatrix}, \quad (4.2)$$

where V is the volume and Br is the residual magnetic flux of the permanent magnet, expressed in Tesla units.




FIGURE 4.1: Diagram of the coordinate system used in the models

4.1.1.1 Electromagnet Dipole Equivalent

Our designs employ electromagnets to generate haptic feedback. We extend the dipole model to include cylindrical and spherical electromagnets used in different devices.

The magnetic field generated by an electromagnet dipole at a point influenced by another dipole or sensor is given by:

$$\mathbf{B}_e(\mathbf{r}, \mathbf{m}_e) = \mathbf{B}(\mathbf{r}, \mathbf{m}_e), \quad (4.3)$$

where \mathbf{r} represents the vector from \mathbf{m}_e to either \mathbf{m}_p or \mathbf{s}_i , as shown in Figure 4.1.

CYLINDRICAL *MagPen* uses a cylindrical electromagnetic. Hence we need a dipole model for cylindrical electromagnets used in our applications. The magnetic field \mathbf{B}_e is modeled using a cylindrical coordinate system aligned with the magnet. The magnet's magnetic moment, \mathbf{m}_e , can be expressed as:

$$\mathbf{m}_e = \alpha m_e [0, 0, 1]^T, \quad (4.4)$$

where α represents the normalized input to the electromagnet and m_e the calibrated strength of the magnet.

SPHERICAL *Omni-v1* and *Omni-v2* consist of custom spherical electromagnets. These operate under different constraints compared to cylindrical electromagnets. However, we can approximate them as the superposition of three orthogonal cylindrical electromagnets. However, in practice the coils are not perfectly orthogonal.

Hence we use a calibration matrix \mathbf{C} (cf. a scalar α in a cylindrical EM). Thus, the magnetic moment of these electromagnets is proportional to the actuation current:

$$\mathbf{m}_e = \mathbf{C} * \mathbf{I}^T, \quad (4.5)$$

where \mathbf{I} is the current vector for each coil, and \mathbf{C} is a calibration matrix derived from measurements.

4.1.1.2 Magnetic Field at a Sensor

To compute the magnetic field at each sensor location, we consider the combined influence of the permanent magnet, the electromagnet, and any background interference:

$$\mathbf{B}_i = \underbrace{\mathbf{B}(\mathbf{r}_{s_i} - \mathbf{r}_p, \mathbf{m}_p)}_{\mathbf{B}_p} + \underbrace{\mathbf{B}(\mathbf{r}_{s_i}, \mathbf{m}_e)}_{\mathbf{B}_e} + \mathbf{B}_n. \quad (4.6)$$

This equation allows us to predict the total magnetic field experienced at each sensor location accurately.

4.1.2 Dipole-dipole Model

We have introduced the dipole model to calculate the magnetic fields produced by both permanent and electromagnets.

Primarily, our objective is to determine the actuation force on the pen, \mathbf{F}_p , which can be computed by integrating the gradient of the magnetic potential over the volume of the pen's permanent magnet:

$$\mathbf{F}_p = \iiint \nabla (\mathbf{M}_p \cdot \mathbf{B}_m(\cdot)) dx dy dz, \quad (4.7)$$

The dipole-dipole interaction model, as described analytically by Yung *et al.* [285], provides an equation for the force between two magnetic dipoles. This model is crucial for determining the necessary magnetic dipole moment of the electromagnet, given the magnetic dipole and position of the permanent magnet to achieve a specific force:

$$\mathbf{F}_p = \frac{3\mu_0}{4\pi r_p^5} \left[(\langle \mathbf{m}_p, \mathbf{r}_p \rangle) \mathbf{m}_e + (\langle \mathbf{m}_e, \mathbf{r}_p \rangle) \mathbf{m}_p + (\langle \mathbf{m}_p, \mathbf{m}_e \rangle) \mathbf{r}_p - \frac{5(\langle \mathbf{m}_p, \mathbf{r}_p \rangle)(\langle \mathbf{m}_e, \mathbf{r}_p \rangle)}{r_p^2} \mathbf{r}_p \right], \quad (4.8)$$

A detailed derivation for specifically planar forces, as used in Optimal Control for Electromagnetic Haptic Guidance Systems, is provided in Appendix A.




FIGURE 4.2: An overview of the control problem. An agent/system takes action \mathbf{a} or input \mathbf{u} based on the state \mathbf{x} or \mathbf{s} . The action/input updates the state. Furthermore the agent receives a reward r .

4.2 OPTIMAL CONTROL

This chapter offers an overview to introduce the reader to the problem formulations and methods of optimal control for discrete-time dynamic systems, which are the main formulations and methods used in the second half of this thesis chapters. This discussion is partially based on Gebhardt’s earlier work and we refer to their publication for an overview [81]. We start by discussing the optimal control problem (OCP) for scenarios where the dynamic system’s model is known. We proceed to examine model predictive control, an optimization technique employed to address these scenarios. Next, we introduce Markov decision processes (MDPs), which serve to frame OCPs when the system model is not known. We extend MDPs to their partial-observable counterparts (POMDPs), these enable scenarios where the system does not have access to the full state space. From there, we explore reinforcement learning (RL), methodologies applied to solve (PO)MDPs. Finally, we review multi-agent reinforcement learning (MARL) for scenarios with multiple agents in a shared environment.

4.2.1 *Optimal Control*

Control theory is a field that studies how to influence the behavior of dynamic systems using external inputs, with the goal of achieving a target state. Optimal control, a subfield of control theory, involves finding the best possible control inputs and state trajectories for a dynamic system. This is done by optimizing a measure of performance or cost. Optimal control can be applied to systems with either continuous or discrete dynamics. In this context, we concentrate on discrete dynamic systems.

First, we introduce the OCP formulation for the case of a discrete system model which inputs are optimized according to a cost function (Figure 4.2). Therefore, we consider the following discrete-time linear time-invariant system

$$\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t), \quad (4.9)$$

where $\mathbf{x} \in \mathbb{R}^n$ is the state vector, $\mathbf{u} \in \mathbb{R}^m$ the input vector, $\mathbf{A} \in \mathbb{R}^{n \times n}$ the system matrix, $\mathbf{B} \in \mathbb{R}^{n \times m}$ the input matrix, $t \in \mathbb{N}_0$ the discrete time, and the pair (\mathbf{A}, \mathbf{B}) is assumed stabilized.

We define $\mathbf{x}(t)$ as the state vector measured at discrete time t and \mathbf{x}_{t+k} be the state vector predicted at discrete time $t+k$ using state equation 4.9 with initial condition $\mathbf{x}_t = \mathbf{x}(t)$. We constraint the problem with:

$$\mathbf{x}(t) \in \mathbb{X} \subseteq \mathbb{R}^n, \mathbf{u}(t) \in \mathbb{U} \subseteq \mathbb{R}^m, \quad (4.10)$$

where \mathbb{X} and \mathbb{U} are the sets of all possible states and inputs respectively.

Then, the a general cost function can be defined as:

$$v_N(\mathbf{x}(t)) = r_T(\mathbf{x}_{t+N}, \mathbf{u}_{t+N}) + \sum_{k=0}^{N-1} \gamma^k r(\mathbf{x}_{t+k}, \mathbf{u}_{t+k}), \quad (4.11)$$

with the stage cost:

$$r(\mathbf{x}_{t+k}, \mathbf{u}_{t+k}) = \mathbf{x}_{t+k}^T \mathbf{Q} \mathbf{x}_{t+k} + \mathbf{u}_{t+k}^T \mathbf{R} \mathbf{u}_{t+k}, \quad (4.12)$$

and terminal cost:

$$r_T(\mathbf{x}_{t+N}, \mathbf{u}_{t+N}) = \mathbf{x}_{t+N}^T \mathbf{P} \mathbf{x}_{t+N}, \quad (4.13)$$

where N is the prediction horizon, $\gamma \in (0, 1]$ a discount factor to give less weight to errors and inputs further in the future, $\mathbf{Q} \in \mathbb{R}^{n \times n}$ and $\mathbf{R} \in \mathbb{R}^{m \times m}$, and $\mathbf{P} \in \mathbb{R}^{n \times n}$ are the state, input and terminal weighting matrices respectively. The weighting matrices are required to be symmetric and positive definite to ensure the objective function is convex. The terminal cost matrix \mathbf{P} is obtained by solving an algebraic Riccati equation, which arises from the infinite-horizon linear quadratic regulator problem [34].

Optimal control problems can be categorized as either finite-horizon (FHOCs) or infinite-horizon (IHOCs) based on the time period considered for optimization. In FHOCs, the control inputs are optimized over a fixed, finite number of future states or time steps. The goal is to find the best control sequence that minimizes the objective function from the initial state to a specified terminal state. We will use a FHOC for haptic control. In IHOCs, the control inputs are optimized assuming the system will continue to operate indefinitely. The objective is to find a control

policy that minimizes the cost over the entire unbounded future, often expressed as an infinite sum of stage costs. The optimal control law is time-invariant - it provides the optimal input as a function of the current state, regardless of the time index. We will use an infinite horizon in our work on Multi-Agent Reinforcement Learning for Point-and-Click Adaptive User Interfaces.

The FHOCP is defined as:

$$v_N^*(\mathbf{x}(t)) = \min_{\mathbf{u}(t)} r_T(\mathbf{x}_{t+N}, \mathbf{u}_{t+N}) + \sum_{k=0}^{N-1} \gamma^k r(\mathbf{x}_{t+k}, \mathbf{u}_{t+k}) \quad (4.14)$$

subject to

$$\mathbf{x}_{t+k+1} = \mathbf{A}\mathbf{x}_{t+k} + \mathbf{B}\mathbf{u}_{t+k}, k = 0, \dots, N-1$$

$$\mathbf{x}_{t+k} \in \mathbb{X}, k = 1, \dots, N$$

$$\mathbf{u}_{t+k} \in \mathbb{U}, k = 0, \dots, N-1$$

$$\mathbf{x}_t = \mathbf{x}(t)$$

with input sequence $\mathbf{u}(t) = (\mathbf{u}_t^T, \dots, \mathbf{u}_{t+N-1}^T)^T \in \mathbb{R}^{Nm}$.

Similarly, the IHOCP is defined as

$$V_\infty^*(\mathbf{x}(t)) = \min_{\mathbf{u}(t)} \sum_{k=0}^{\infty} \gamma^k r(\mathbf{x}_{t+k}, \mathbf{u}_{t+k}) \quad (4.15)$$

$$\text{subject to } \mathbf{x}_{t+k+1} = \mathbf{A}\mathbf{x}_{t+k} + \mathbf{B}\mathbf{u}_{t+k}, k = 0, 1, \dots \quad (4.16)$$

$$\mathbf{x}_{t+k} \in \mathbb{X}, k = 1, 2, \dots$$

$$\mathbf{u}_{t+k} \in \mathbb{U}, k = 0, 1, \dots$$

$$\mathbf{x}_t = \mathbf{x}(t)$$

with input sequence $\mathbf{u}(t) = (\mathbf{u}_t^T, \mathbf{u}_{t+1}^T, \dots)^T$.

We introduced the OCP for a linear system so far. However, not all systems are linear or can be linearized; such as human behavior. In the nonlinear case, the right part of (4.9) is substituted by a nonlinear function that describes the evolution of the dynamic system:

$$\mathbf{x}(t+1) = F(\mathbf{x}(t), \mathbf{u}(t)). \quad (4.17)$$

The terminal penalty for non-linear system is given by the Lyapunov equation [115], rather than the Riccati equation.

4.2.2 Model Predictive Control

Model predictive control (MPC) is a practical implementation of (in)finite horizon optimal control for the measured state vector $\mathbf{x}(t)$ to attain the predicted optimal

input sequence $\mathbf{u}^{(t)}$. Specifically, it is a closed-loop control method that measures the state (\mathbf{x}) at every time step (t), solves the optimal control equations (4.14), and then applies its first element ($\mathbf{u}^{(t)}$) to the system. This is repeated at each discrete time t with a receding prediction horizon. Thus, MPC is also denoted as receding horizon control (RHC).

Both the FHOCP and the IHOCP can be reformulated as a quadratic programming (QP) problem for linear systems with quadratic cost functions and linear constraints ([90]):

$$\begin{aligned} & \underset{\mathbf{X}}{\text{minimize}} \quad \frac{1}{2} \mathbf{X}^T \mathbf{H} \mathbf{X} + \mathbf{f}^T \mathbf{X} & (4.18) \\ & \text{subject to} \quad \mathbf{A}_{ineq} \mathbf{x} \leq \mathbf{b}_{ineq} \\ & \quad \text{and} \quad \mathbf{A}_{eq} \mathbf{x} = \mathbf{b}_{eq} , \end{aligned}$$

Here, \mathbf{X} represents the stacked state vectors \mathbf{x}_t and inputs \mathbf{u}_t at each time point. The matrices \mathbf{H} and \mathbf{f} contain the quadratic and linear cost coefficients, respectively, as defined in (4.11). The matrices \mathbf{A}_{ineq} and \mathbf{b}_{ineq} describe the linear inequality constraints on states and inputs from (4.10), while \mathbf{A}_{eq} and \mathbf{b}_{eq} represent the linear equality constraints from our model in (4.9) for each time point $k \in 0, \dots, N$.

To solve nonlinear dynamic systems and cost functions, we can use nonlinear programming (NLP) or sequential quadratic programming (SQP) methods. NLP addresses the nonlinear optimization problem directly, whereas SQP iteratively solves local quadratic programming (QP) approximations of the nonlinear problem. The selection of horizon length, terminal cost, and constraints in these formulations is crucial for ensuring closed-loop stability and performance. We use numerical solvers to address MPC problems [63].

4.2.3 Markov Decision Processes

Model Predictive Control (MPC) addresses the OCP of a system assumed to behave deterministically according to a known dynamic model. However, many OCPs inherently have stochastic state transitions, making it impossible to model the underlying system dynamics accurately. Markov Decision Processes (MDPs) offer a mathematical framework for these situations, where a system model and suitable control policies are learned from the system's interactions with its environment. Generally, MDPs can be used to represent a wide range of sequential decision-making problems in stochastic settings [121].

An MDP satisfies the Markov property, which states that state transitions depend only on the current state and action, not on any prior states or actions. An MDP is

defined by a five-tuple (S, A, T, R, γ) , where S is the set of states, A is the set of actions, $T : S \times A \times S \rightarrow [0, 1]$ is the transition probability function, and $T(\mathbf{s}', \mathbf{a}, \mathbf{s})$ represents the probability of transitioning from state \mathbf{s}' to state \mathbf{s} after taking action \mathbf{a} . $R : S \times A \rightarrow \mathbb{R}$ is the reward function, with rewards discounted by a factor γ . The expected discounted reward for taking action \mathbf{a} in state \mathbf{s} under policy π is called the Q value:

$$Q^\pi(\mathbf{s}, \mathbf{a}) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t R(\mathbf{s}_t, \mathbf{a}_t) \mid \mathbf{s}_t = \mathbf{s}, \mathbf{a}_t = \mathbf{a} \right], \quad (4.19)$$

The value function $V^\pi(s)$ is the expected return starting from state s and following policy π :

$$V^\pi(s) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t R(\mathbf{s}_t, \mathbf{a}_t) \mid \mathbf{s}_t = \mathbf{s} \right] \quad (4.20)$$

The Q values and value function are related via:

$$V^\pi(\mathbf{s}) = \sum_{\mathbf{a}} \pi(\mathbf{a} \mid \mathbf{s}) Q^\pi(\mathbf{s}, \mathbf{a}) \quad (4.21)$$

The Q values of following states are related via the Bellman equation:

$$Q^\pi(\mathbf{s}, \mathbf{a}) = \sum_{\mathbf{s}'} P(\mathbf{s}' \mid \mathbf{s}, \mathbf{a}) [R(\mathbf{s}', \mathbf{s}, \mathbf{a}) + \gamma Q^\pi(\mathbf{s}', \pi(\mathbf{s}'))]. \quad (4.22)$$

The optimal policy can then be computed as $\pi^* = \arg \max_{\mathbf{a}} Q^\pi(\mathbf{s}, \mathbf{a})$.

4.2.4 Partially Observable Markov Decision Processes

Partially Observable Markov Decision Processes (POMDP) extend the MDP, and is a mathematical framework for single-agent decision-making in stochastic partially observable environments [11], which is a generalization over Markov Decision Processes [103]. A POMDP is a seven-tuple $(S, O, A, T, F, R, \gamma)$. In POMDPs, the exact states ($s \in S$) of the evolving environment may or may not be captured fully. Therefore, observations ($\mathbf{o} \in O$) represent the observable states, which may differ from the exact state. Similar to T , $F : S \times A \times O \rightarrow [0, 1]$ is an observation probability function, where $F(\mathbf{o}, \mathbf{a}, \mathbf{s}')$ is the probability of observing \mathbf{o} while transitioning to \mathbf{s}' after taking action \mathbf{a} .

The Q-value function for a POMDP is defined over belief states and actions, rather than exact states and actions. Let b denote a belief state, which is a probability distribution over states, so that $b'(\mathbf{s}')$ is the updated belief of being in state \mathbf{s}' . To compute the belief state in a Partially Observable Markov Decision Process

(POMDP), we use a recursive state estimation process based on the previous belief state, the action taken, and the current observation received. Then, the belief state update equation is:

$$b'(s') = \eta F(\mathbf{o}|s', \mathbf{a}) \sum_{\mathbf{s} \in \mathcal{S}} T(s'|\mathbf{s}, \mathbf{a}) b(\mathbf{s}) \quad (4.23)$$

where $\eta = 1/P(\mathbf{o}|b, \mathbf{a})$ is a normalizing constant with:

$$P(\mathbf{o}|b, \mathbf{a}) = \sum_{s' \in \mathcal{S}} F(\mathbf{o}, s', \mathbf{a}) \sum_{\mathbf{s} \in \mathcal{S}} T(s'|\mathbf{s}, \mathbf{a}) b(\mathbf{s}). \quad (4.24)$$

This process is repeated after each action and observation to maintain a current belief state that summarizes all information received so far. The belief state forms a sufficient statistic for the entire action-observation history and allows the POMDP to be cast as a continuous-state MDP, called a belief MDP

The value function for a POMDP policy π over belief states is defined as:

$$V^\pi(b) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t R(\mathbf{s}_t, \mathbf{a}_t) \mid b_0 = b \right] \quad (4.25)$$

Then the Q-value function for a POMDP policy π (casted as belief MDP) is defined as:

$$Q^\pi(b, \mathbf{a}) = \sum_{\mathbf{s} \in \mathcal{S}} b(\mathbf{s}) \left[R(\mathbf{s}, \mathbf{a}) + \gamma \sum_{\mathbf{o} \in \mathcal{O}} F(\mathbf{o}, \mathbf{s}, \mathbf{a}) V^\pi(b_{\mathbf{o}}^{\mathbf{a}}) \right], \quad (4.26)$$

where $b_{\mathbf{o}}^{\mathbf{a}}$ is the updated belief state after taking action \mathbf{a} and observing \mathbf{o} and $V^\pi(b_{\mathbf{o}}^{\mathbf{a}})$ is the value of the updated belief state $b_{\mathbf{o}}^{\mathbf{a}}$ under policy π . The optimal policy can then be computed as $\pi^* = \arg \max_{\mathbf{a}} Q^\pi(b, \mathbf{a})$

4.2.5 Stochastic Games

MDPs or POMDPs assume a single policy. Stochastic games generalize MDPs for multiple policies [237]. When players do not have perfect information about the environment, stochastic games become partially observable stochastic games. A partially observable stochastic game is defined as an eight-tuple $(N, \mathcal{S}, \mathcal{O}, \mathcal{A}, T, \mathcal{F}, \mathcal{R}, \gamma)$, where N is the number of policies. $\mathcal{S} = S_1 \times \dots \times S_N$ is a finite set of state sets, and $\mathcal{O} = O_1 \times \dots \times O_N$ is a finite set of observation sets, with subscripts indicating different policies. $\mathcal{A} = A_1 \times \dots \times A_N$ is a finite set of action sets. $\mathcal{F} = F_1 \times \dots \times F_N$ defines a set of observation probability functions of different players. A set of reward functions is defined as $\mathcal{R} = R_1, \dots, R_N$. Furthermore, we define a set of policies as $\Pi = \pi_1, \dots, \pi_N$.

All policies have their individual actions, states, observations, and rewards. In this paper, we optimize each policy individually, while the observations are influenced by each other’s actions, hence we can treat this as multiple distinct POMDPs as outlined in Section 4.2.4.

Interactive POMDPs (I-POMDPs) are an alternative framework for multi-agent decision making [89]. A core different with partially observable stochastic games is the I-POMDPs explicitly model belief states, and take actions base on them. Whereas stochastic games involve less cognitive modeling and have an implicit representation.

4.2.6 Reinforcement Learning

Reinforcement learning (RL) solves MDPs by learning a state-action value function $Q(s, a)$ (or $Q(b, a)$) that approximates the Q value defined by the Bellman equation $Q^\pi(s, a)$. RL algorithms fall into two categories: model-based and model-free. In model-based algorithms, state transition probabilities are known, and policies are determined by enumerating possible state sequences following an initial state and action, summing the expected rewards along these sequences. This dissertation focuses on model-free RL algorithms for solving MDPs, so for an overview of model-based approaches, see [184, 213]. In model-free RL, the transition probability functions are unknown but can be sampled. These algorithms learn the approximate state-action value function $Q(s, a)$. In Deep RL, the policy π_θ is learned as a multi-layered perceptron (MLP), where θ represents the learnable parameters, or the weights of the MLP.

Trust Region Policy Optimization (TRPO) [230] is a policy optimization algorithm designed to achieve stable and efficient training in reinforcement learning by maintaining a trust region around the current policy. The algorithm optimizes the policy by solving the following constrained optimization problem:

$$\arg \max_{\theta} \mathbb{E}_{\mathbf{s}_t, \mathbf{a}_t \sim \pi_{\theta_{\text{old}}}} \left[\frac{\pi_{\theta}(\mathbf{a}_t | \mathbf{s}_t)}{\pi_{\theta_{\text{old}}}(\mathbf{a}_t | \mathbf{s}_t)} A^{\pi_{\theta_{\text{old}}}}(\mathbf{s}_t, \mathbf{a}_t) \right], \quad (4.27)$$

subject to a constraint on the Kullback-Leibler (KL) divergence between the new and old policies:

$$\mathbb{E}_{\mathbf{s}_t \sim \rho_{\theta_{\text{old}}}} [D_{\text{KL}}(\pi_{\theta_{\text{old}}}(\cdot | \mathbf{s}_t) | \pi_{\theta}(\cdot | \mathbf{s}_t))] \leq \delta, \quad (4.28)$$

where π_θ represents the policy parameterized by θ , and $A^{\pi_{\theta_{\text{old}}}}(\mathbf{s}_t, \mathbf{a}_t)$ is the advantage function estimated under the old policy $\pi_{\theta_{\text{old}}}$:

$$A^\pi(\mathbf{s}, \mathbf{a}) = Q^\pi(\mathbf{s}, \mathbf{a}) - V^\pi(\mathbf{s}), \quad (4.29)$$

The constraint ensures that each policy update is conservative, limiting the divergence between successive policies, thereby enhancing the stability and performance of the learning process. This approach mitigates the risk of drastic policy changes, which can lead to performance degradation or instability in the learning process.

In this dissertation we use Proximal Policy Optimization (PPO) [231], which is a simplified variant of TRPO designed to retain its benefits while reducing computational complexity. PPO maximizes a clipped objective function:

$$\arg \max_{\theta} \mathbb{E}_t [\min(r_t(\theta) \mathbf{a}_t, \text{clip}(r_t(\theta), 1 - \varepsilon, 1 + \varepsilon) \mathbf{a}_t)], \quad (4.30)$$

where $r_t(\theta) = \frac{\pi_{\theta}(\mathbf{a}_t | s_t)}{\pi_{\theta_{\text{old}}}(\mathbf{a}_t | s_t)}$ is the probability ratio between the new and old policies, and \mathbf{a}_t is the estimated advantage function. The clipping function constrains $r_t(\theta)$ within a range of $[1 - \varepsilon, 1 + \varepsilon]$, with ε being a small hyperparameter. This restriction prevents excessively large updates that could destabilize training, striking a balance between exploration and exploitation. PPO's design simplifies the implementation of policy updates by eliminating the need for complex constraints on KL divergence, thus making it a practical and widely adopted algorithm in reinforcement learning applications.

Part II

THE DESIGN OF SHARED VARIABLE
INTERFACES

CONTACT-FREE NON-PLANAR HAPTICS WITH A SPHERICAL ELECTROMAGNET

We investigated a shared variable interface using a contact-free volumetric haptic feedback device. This device employs a symmetric electromagnet in combination with a dipole magnet model and a simple control law to deliver dynamically adjustable forces onto a handheld tool. The tool requires only an embedded permanent magnet, allowing it to remain completely untethered. Despite being contact-free, the force is grounded via the spherical electromagnet, enabling the user to feel relatively large forces (up to 1N at contact). The device can render both attracting and repulsive forces within a thin shell around the electromagnet. We conducted a user experiment with six participants to characterize the force delivery aspects and perceived precision of our system. Our findings indicate that users can discern at least 25 locations for repulsive forces. This research marks the first step towards understanding shared variable interfaces. By eliciting forces on the tool manipulated by the user, we enable haptic feedback in the same modality as the user's exertion, resulting in natural and intuitive interactions.

5.1 INTRODUCTION

A specific instance of a shared variable interface is a subset of haptic interfaces. Haptic interfaces are particularly interesting because they operate in the physical world, where the challenge of variable ownership is more impactful. Additionally, two agents exerting force on a single object intuitively exemplify shared variable interfaces. Furthermore, many emerging computing paradigms, such as virtual and augmented reality (VR/AR), rely on haptic feedback as an additional information channel to enhance the user experience. For example, in VR, haptic feedback increases the sense of presence and immersion by rendering collisions, shapes, and forces between the user and virtual objects.

For a haptic device to qualify as having a shared variable interface, it must function as both an input and output device. Many haptic devices primarily function as output devices, such as vibrotactile actuators embedded in handheld controllers [274] or worn on the body [99]. Vibrotactile feedback offers one-way communication and can only render coarse, non-localized haptic sensations. In contrast, our device allows for bi-directional communication on a shared variable. While complex setups such as articulated arms and exoskeletons can render large-force haptic feedback in three-dimensional space, they typically require force anchoring in the environment and involve cumbersome, bulky mechanisms, which hinder user uptake in walk-up-and-use scenarios.

To address the limitations of vibrotactile and large complex devices, we propose a contact-free, volumetric haptic feedback approach via an omnidirectional electromagnet. The device consists of a single 60 mm diameter spherical electromagnet capable of rendering attractive and repulsive forces onto permanent magnets embedded in pointing tools, such as a stylus or magnets worn on the user's fingertip. Leveraging a dipole-dipole approximation of the electromagnet-magnet interaction, our system can calculate and control the forces exerted on the permanent magnet in real-time, dynamically adjusting the force perceived by the user. The system can deliver perceptible forces up to 1N in a thin volume above the surface. Furthermore, we demonstrate that users can distinguish at least 25 different set-points separated by 18° on the surface of the sphere.

To demonstrate the efficacy of our approach, we designed a functional prototype comprising an iron core and three custom-wound copper coils. The electromagnet is encased in a plastic dome upon which tools can come into contact and move about its surface (see Figure 5.1). The prototypical system can render radial (along the vector from the magnet to the tool) and tangential forces, both in attractive and repulsive polarities. The system can dynamically adjust the opening angle and steepness of




FIGURE 5.1: We introduce a novel contact-free mechanism to render haptic feedback onto a tracked stylus via a hemispherical electromagnet. An approximate model of the magnet interaction and a computationally efficient control strategy allow for the dynamic rendering of attracting and repulsive forces, for example, allowing users to explore virtual surfaces in a thin shell surrounding the device (inset).

the electromagnetic potential to gently guide the user towards a desired set-point in the thin volume above the device.

Modulating the magnetic field based on tool position opens the door to various interactive applications. For example, in virtual terrain exploration, the tool can be repelled when moved along mountains and attracted to valleys while descending (see Figure 5.1, inset). Another example is creating the sensation of stirring a viscous liquid by emulating the fluid’s drag on the tool. To enable these interactive experiences, our device builds on three key components representing our contributions in this chapter:

- A computational model based on magnetic dipole-dipole interaction to produce force maps that allow for designing and generating location-dependent feedback,

- The design and implementation of a 3 degree-of-freedom (DoF) spherical electromagnet prototype,
- A control strategy that translates desired high-level forces into low-level input signals (currents/voltages) for the coils, fast enough for interactive use.

To assess the efficacy of the proposed design, we experimentally characterized the system properties and conducted a perceptual study exploring the thresholds for perception and localization capabilities of the electromagnetic actuation approach. Results from these early user tests indicate that users can perceive at least 25 different spatial locations with high precision.

5.2 SYSTEM DESCRIPTION

We introduce a haptic feedback system that enables dynamic interactions with virtual surfaces through an untethered, contact-free tool. Our device is a hemispherical shell. The core consists of three coils with mutually orthogonal axes. By controlling the current flow through the coils, we can shape the magnetic field around the device. This, in turn, enables the device to exert controlled electromagnetic forces on the permanent magnet located inside a handheld tool such as a stylus. Despite being contact-free, the forces perceived by the user are ultimately grounded to the support onto which the device is mounted, allowing for comparatively strong feedback. We now detail the main components that make up our contribution: 1) a computational model of the electromagnet-magnet interactions; 2) the prototypical hardware design; and 3) a real-time control algorithm.

5.2.1 *Haptic force mapping*

To enable the envisioned interactive experiences, we must be able to dynamically adjust the haptic feedback. We therefore require a model for the magnetic interaction between device and tool that is 1) precise enough to predict forces with sufficient accuracy and 2) fast enough to run at the feedback rates required for haptic interaction.

Computing the magnetic field around, and resulting interaction between, arbitrarily-shaped objects is a challenging and computationally expensive task. However, even though the magnetic field can be very complex in the direct vicinity of an object, this complexity rapidly decays with increasing distance and approaches a simple dipole field. This fact has been exploited in previous work to construct fast, approximate models based on dipole-dipole interaction [264]. Instead of solving the Maxwell equations on a discretization of ambient space, this approximate model only requires

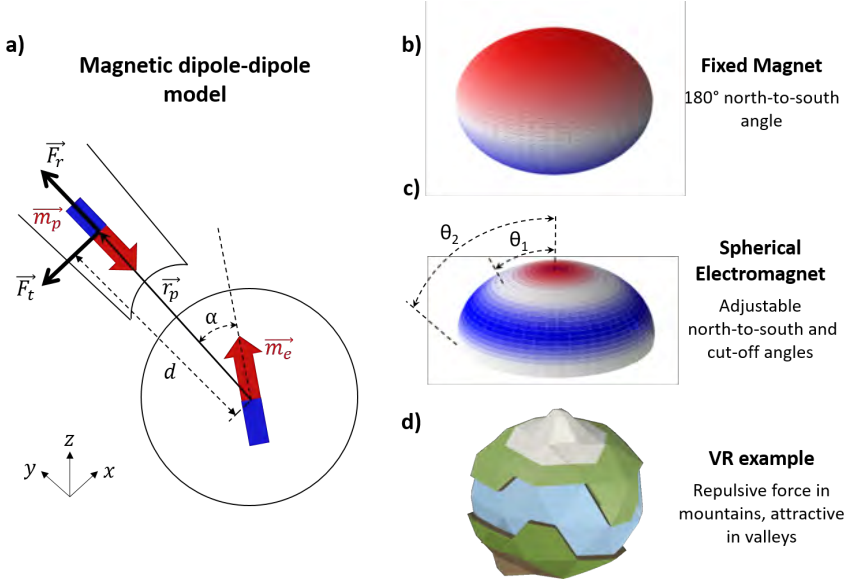


FIGURE 5.2: Schematic of the main quantities necessary to compute desired radial and tangential forces (a). Insets show: force map of a permanent magnet (b). Adjustable force map generated by our approach (c). Here $\mathbf{r}_0 = d_{min}\mathbf{e}_z$, $\theta_1 = \pi/10$ and $\theta_2 = 3\pi/10$. Example virtual surface that can be felt by the user (d).

the magnitude and orientation of the magnetic moment of each dipole, leading to drastically reduced computation times.

In adopting this approach, we model both the electromagnet of the device and the tool as a single dipole (see Figure 5.2.a). Let $\mathbf{m}_p, \mathbf{m}_e \in \mathbb{R}^3$ denote the magnetic moments of the permanent magnet in the tool and the electromagnet in the device, respectively. The forces exerted on the tool, expressed in local coordinates, are obtained as:

$$\mathbf{F}_r = -\frac{3\mu_0 m_e m_p}{2\pi d^4} \cos(\alpha) \mathbf{e}_r, \quad (5.1)$$

$$\mathbf{F}_t = -\frac{3\mu_0 m_e m_p}{4\pi d^4} \sin(\alpha) \mathbf{e}_t, \quad (5.2)$$

where $m_e = |\mathbf{m}_e|$, $m_p = |\mathbf{m}_p|$. In the above expression, \mathbf{F}_r is the force in the radial direction $\mathbf{r}_p = d \mathbf{e}_r$ from the center of the device to the tool. Likewise, \mathbf{F}_t is the force in the tangential direction \mathbf{e}_t that tends to align the location of the two dipoles along \mathbf{e}_r . Assuming that the tool is in contact with the shell, both force components

depend only on the relative angle α between the dipoles. Furthermore, \mathbf{F}_r and \mathbf{F}_t are attractive (negative) when the two dipoles have the same sign and $\alpha < \pi/2$. Conversely, the forces become repulsive (positive) when the dipoles have opposite orientations (see Figure 5.2.a).

The interaction forces decay quickly, as $1/d^4$, with increasing magnet-magnet distance. The maximum force $\mathbf{F}_{r,\max}$ is obtained when the tool is in contact with the device ($d = d_{\min}$). In our case, $d_{\min} = 50$ mm, since the outer case radius is 30 mm and inside the tool, the magnet center is 20 mm away from the tool tip. Our proposed geometry ensures that the distance d will remain constant across the working surface as long as the tool is kept in contact with the surface, allowing for much simpler control of the force. However, it is worth noting that moving the tool 1 cm away in the radial direction makes the force fall to approximately $\mathbf{F}_{r,\max}/2$, another extra centimeter results in a force $\mathbf{F}_{r,\max}/4$. This rapid decay of the interaction forces can, to some extent, be mitigated by increasing the intensity of the magnetic field. However, to maintain power consumption and thermal effects within reasonable bounds, we constrain our interactions to a volumetric shell ($d_{\min} \leq d \leq d_{\min} + 2\text{cm}$) above the device's surface.

Equations 5.1 and 5.2 also reveal the comparatively weak variation of force magnitude with respect to angle that one would expect when two magnets interact: switching from attractive to repulsive forces requires a change in orientation of $\alpha = \pi$; see Figure 5.2.b. This weak force variation is inherent to permanent magnets: whereas the far-field interaction is dominated by torque (which decays only as $1/d^3$), the near-field force interaction is governed by the location of the dipoles, not their orientation. In our setting, this property would translate into weak angular resolution with a permanent magnet. To address this problem, we introduce the concept of a *force map* that uses magnetic pole transformation to take advantage of the spherical symmetry and that is compliant with the physics of the system. Our system can generate force maps equivalent to multiple alternating pole regions, having sharper repulsive domes and attractive valleys. The force map is defined by four parameters:

- The center \mathbf{r}_0 of the potential. When rendering a mountain-like dome, for instance, \mathbf{r}_0 is the summit.
- The height of the dome is measured as the maximum magnetic moment intensity m_{e0} .
- The angle (θ_1) (i.e., the location of the tool in polar coordinates with respect to \mathbf{r}_0) where the radial force vanishes for the first time. In our example, (θ_1) is the angle from the summit to the base.

Algorithm to calculate desired forces

% To compute \mathbf{m}_e given the tool position and the force map.

Require: $\text{calc_Me}(\mathbf{r}_p, \mathbf{r}_0, m_{e0}, \theta_1, \theta_2)$:

$$\mathbf{r}_p|_{\mathbf{r}_0} = \mathbb{T}_{\mathbf{r} \rightarrow \mathbf{r}_0} \cdot \mathbf{r}_p$$

$$\mathbf{F}|_{\mathbf{r}_0} = \text{calc_F}(\mathbf{r}_p|_{\mathbf{r}_0}, \mathbf{r}_0, m_{e0}, \theta_1, \theta_2)$$

$$\mathbf{F} = (\mathbb{T}_{\mathbf{r} \rightarrow \mathbf{r}_0})^{-1} \cdot \mathbf{F}|_{\mathbf{r}_0}$$

$$\mathbf{F}|_{\mathbf{r}_p} = \mathbb{T}_{\mathbf{r} \rightarrow \mathbf{r}_p} \cdot \mathbf{F}$$

$$\mathbf{m}_e|_{\mathbf{r}_p} = \frac{4\pi d^4}{3\mu_0} [1, 1, -1/2] \cdot \mathbf{F}|_{\mathbf{r}_p}$$

$$\mathbf{m}_e = (\mathbb{T}_{\mathbf{r} \rightarrow \mathbf{r}_p})^{-1} \cdot \mathbf{m}_e|_{\mathbf{r}_p}$$

return \mathbf{m}_e

% To compute the actuation force in the $|_{\mathbf{r}_0}$ coordinates.

Require: $\text{calc_F}(\mathbf{r}_p|_{\mathbf{r}_0}, \mathbf{r}_0, m_{e0}, \theta_1, \theta_2)$:

$$F_r = 0$$

$$F_t = 0$$

if $d < d_{max}$ **and** $\alpha < \theta_2$ **then**

$$F_r = 2F_0 \cos(\alpha \frac{2\theta_1}{\pi}) \left(\frac{\|\mathbf{r}_0\|}{\|\mathbf{r}_p\|} \right)^4$$

$$F_t = F_0 \sin(\alpha \frac{2\theta_1}{\pi}) \left(\frac{\|\mathbf{r}_0\|}{\|\mathbf{r}_p\|} \right)^4$$

end if

$$\mathbf{F}|_{\mathbf{r}_0} = [F_r, F_t, 0]$$

return $\mathbf{F}|_{\mathbf{r}_0}$

FIGURE 5.3: Pseudo-code of our force calculation algorithm. Note that $\mathcal{T}_{r_i \rightarrow r_j}$ is the rotation matrix that maps from coordinate system r_i to r_j , and that $\mathbb{T}_{r_j \rightarrow r_i} = (\mathbb{T}_{r_i \rightarrow r_j})^{-1} = (\mathbb{T}_{r_i \rightarrow r_j})^T$.

- The cut-off angle θ_2 after which the potential is set to be zero. Having such a cut-off mechanism allows us to control how many individual potentials can be combined into one force map without mutual interference.

Figure 5.3 summarizes our algorithm to calculate the actuation vector \mathbf{m}_e given the tool position and force map as input. For simplicity and efficiency, we perform the different calculations in their natural coordinate system: the Cartesian system $\mathbf{r} = [x, y, z]$, the spherical system relative to the map's center \mathbf{r}_0 , and the spherical system centered around the tool position \mathbf{r}_p .

The force calculation incorporates the angular scaling by using $(\alpha \frac{2\theta_1}{\pi})$ as an argument for the trigonometric functions in Equations 5.1 and 5.2. Note that if $\theta_1 = \pi/2$, we recover a permanent magnet. In Figure 5.2.c, we show an example where the center of the potential (*red*) is on the north pole of the sphere, the first vanishing region (*white*) appears at 18° and the forces are cut off at 54° (*blue*).

Using the algorithm described in Figure 5.3, we obtain at each time step an actuation input $\mathbf{m}_e = (m_{e-x}, m_{e-y}, m_{e-z})^T$ given the tool position. Depending on the requirements of the application, the potential parameters (center position, intensity, angular variation, and cut-off) may also change as a function of tool position. For example, the force map for the terrain example can be dynamically adapted to emulate changes in landscape over time.

5.2.2 Spherical electromagnetic actuator

Having laid out the computational model for generating haptic feedback based on dipole interactions, we now describe hardware and implementation aspects for rendering these forces on our device (Fig. 5.1). Our device renders haptic forces by controlling the magnetic field generated by a spherical electromagnet. Compared to other alternatives, this approach has several advantages. First, there are no mechanically moving parts in the actuator, reducing complexity and eliminating wear. Changing the orientation of the resulting force on the tool is accomplished by adapting the currents in each coil such as to rotate the induced dipole in the core as desired; see also Figure 5.4. The underlying physical principle is that, in the presence of linear and isotropic materials, the magnetic field $\mathbf{B}(\mathbf{r})$ at any given point \mathbf{r} can be calculated as the sum over all contributions of all magnetic sources [211]. Under this linearity property of \mathbf{B} , the magnetic field produced by the three orthogonal coils is the superposition of the fields generated by each coil individually. Finally, we insert a magnetic core with isotropic (i.e., spherical) geometry and material at the center of the coils and operate it in the linear regime (i.e. $m_e \ll m_{saturation}$), linearity is maintained such that $\mathbf{B}(\mathbf{r})$ can be computed by summing up each coil's contributions.

In order for the previous statement to remain valid, two assumptions have to be made. First, hysteresis effects can be neglected: the lower the coercivity and remanence of the core material, the lower the effect of past states of the electromagnet on the current one. The second assumption is that the distance d between dipoles is large enough such that the core magnetization due to m_p is small compared with the effect of the coils. This will not be true if, for example, the tool snaps to the sphere with no electrical current in the coils. In our setting, however, a spherical cap around




FIGURE 5.4: 3D cross-section of the proposed hardware setup. The device measures 15×15 cm across the base. Three coils are placed, orthogonal to each other and surrounding the iron core. Forces can be rendered onto a permanent magnet moving above the device. Hall sensors are used for calibration. A plastic cover isolates the coils from the user thermally and electrically. Active cooling is provided via several fans mounted in the base.

the coils prevents too close approach of the tool and, at the same time, provides the grounding required for generating sufficiently strong interaction forces.

For the standard low-carbon steel core, we did not observe any hysteresis effects for the update of \mathbf{m}_e at 50 Hz refresh rate. To avoid the undesired self-magnetization of the core due to the tool, we tuned the size of the permanent magnet and the coil parameters using FEM simulations, followed by minor design adjustments informed by real-world tests.

The design choices for the hardware of our prototype are motivated by our goal to develop a device that is affordable and easy to manufacture. In particular, we use off-the-shelf electronic components but custom-wound coils. FEM simulations in Comsol Multiphysics are used to assist in the exploration of the design space. In Figure 5.4 we show a 3D CAD rendering of our device. The external dimensions are 150 mm by 150 mm by 95 mm. The structure is built out of laser-cut acrylic glass and 3D-printed parts. The three orthogonal coils are arranged around the 30 mm steel core. All coils have a resistance of roughly 0.6Ω at room temperature. We use the 12V line of a standard CPU power supply to drive the coils, meaning a maximum electrical current of 24A per coil at full strength. The electrical current in each coil is controlled by a high-power motor driver (Pololu 18v17). The PWM signals are generated by a 12-bit driver (PCA9685) that allows for easy tuning of the carrier frequency and the duty cycle with 12-bit resolution. To be able to accurately control

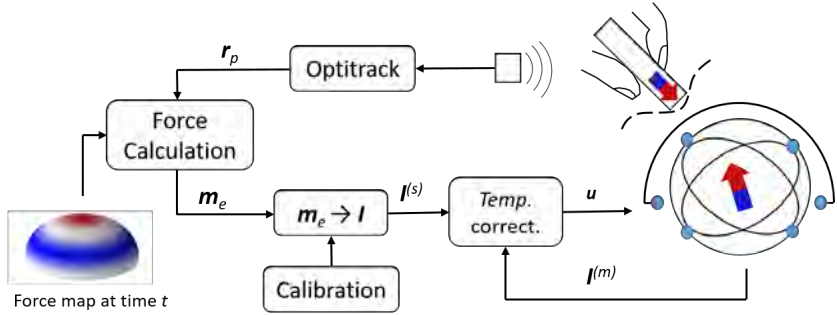


FIGURE 5.5: Schematic overview of the software pipeline. Given the desired force map at time t , and the tool position provided by an external tracking system, we calculate the input value \mathbf{m}_e using the algorithm of Fig. 5.3. Then the system inputs are computed using Eq. 5.4, and finally a temperature compensation step corrects the system inputs.

the electrical current and compensate for thermal drifts, we use INA219 current sensors in each coil with a 0.01Ω shunt resistor. Finally, an Arduino board creates the bridge between the I2C components and the PC. The hardware is completed by 9 Hall sensors arranged collinearly with the axes and diagonals of the coils. Six fan coolers below the coils provide active cooling.

5.2.3 Control Strategy

The main objective of the actuator control loop is to generate a stable and controllable force on the haptic tool. Although the mathematical principles are straightforward, the practical implementation poses some problems. Since the magnetic field is directly proportional to the current (Fig. 5.7), controlling the latter is sufficient to determine the state of the system. If the resistance is known, controlling the voltage is equivalent to controlling the current via Ohm's law,

$$I = V/R, \quad (5.3)$$

and the voltage in turn can be controlled via Pulse-Width Modulation (PWM). Therefore, the input to our system is the PWM frequency. The complete control loop is shown in Figure 8.3. However, significant heating occurs due to the necessary power that in turn increases the resistance. Therefore, the PWM duty cycle (i.e., voltage) needs to be adjusted to maintain a constant current. Measuring the current allows determination of the resistance via inversion of Eq 5.3. A simple controller

then computes an input $\mathbf{u} \in [-1, 1]$ at time t , corresponding to the PWM duty cycle. This depends on the desired current in Ampere ($\mathbf{I}^{(s)}_t$), the resistance in Ohm (\mathbf{R}_t), and the maximum voltage in the system, $V_0 = 12$:

$$\mathbf{u}_t = \frac{\mathbf{I}^{(s)}_t}{V_0} * \mathbf{R}_t \quad , \quad (5.4)$$

where $\mathbf{I}^{(s)}_t$ is based on the desired magnetization, \mathbf{m}_e , computed via the algorithm presented in Fig. 5.3 and can be determined via Biot-Savart Law (adapted for our purpose):

$$\mathbf{I}^{(s)}_t = \mathbf{c} * \frac{\mathbf{m}_e * \mu_0}{2 * \pi * \mathbf{d}^3} \quad , \quad (5.5)$$

here \mathbf{c} is a constant coming from a calibration procedure that, with the help of five Hall sensors, maps input current to \mathbf{m}_e (Fig. 5.7). $\mu_0 = 4 * \pi * 10^{-7}$ is the relative permeability of air and \mathbf{d} is the distance from the core to the Hall sensors used for calibration (0.055 meter). Due to the thermal effects, \mathbf{R}_t however is not a constant but depends on the measured current ($\mathbf{I}^{(m)}$) computed and averaged over a sliding window:

$$\mathbf{R}_t = \frac{V_0 * \frac{1}{N} \sum^N \mathbf{u}_t - i}{\frac{1}{N} \sum^N \mathbf{I}^{(m)}_{t-i}} \quad . \quad (5.6)$$

5.3 SYSTEM EVALUATION

One of the main physical limitations of EM-based systems is thermal effects due to Joule heating, to obtain large forces [67]. The temperature is directly proportional to the actuation power (P) and the thermal dissipation obtained by the active and/or passive cooling. We evaluated the thermal behavior of our system for different power values. In this experiment, we set the current to on' for three minutes and then let the device cool down. Figure 5.6 shows data from the middle coil actuated at PWM = 30%. T_{out} is the temperature measured at the coil boundary, measured with a Dallas DS18B20 sensor. T_{in} is the average temperature of the copper wire obtained via the variation in resistance. We also plot the electrical current I_y that drops as the coil heats up and the resistance increases. Note that no temperature compensation was used for building these thermal calibration curves. Each coil is able to accumulate some heat during the actuation and continuously dissipates it by the forced air circulation. Our system has a thermal time (τ_T) on the order of minutes, in which it reaches the asymptotic temperature. The average power in the past τ_T seconds must be maintained within a safe value P_{ave} . Based on this plot, we choose $P_{ave} = 17W$ per coil for our system. However, each coil can absorb peaks up to $15 * P_{ave}$ for a few seconds.




FIGURE 5.6: Thermal characterization of one of the coils as function of time. During the first 3 minutes the y-coil is driven with PWM=30%, and then we let it cool over the remaining 3 minutes. T_{in} is calculated by taking the thermally caused resistance variations into account while the current I_y is ‘on’, and T_{out} is measured.

Within this safe range, we calibrate the values of m_e for each axis as a function of the current in each coil with the hall sensors around the sphere (see Figure 5.4) and with Eq. 5.5. Figure 5.7 shows the experimentally attained magnetization in the core \mathbf{m}_e as a function of the current. For reference, applying a power $P_0 = 100$ W to each coil ($I_i = 12.9$ A), the equivalent dipole is $\mathbf{m}_e = [2.52; 2.7; 2.82]$ Am². We also obtain non-zero terms away from the diagonal since the coils are not perfectly orthogonal, and we use the calibration data to correct the PWM duty cycles.

Finally, values for the force acting on the permanent magnet can be attained via setting the magnetic dipole of the tool and Eq. 5.1 and 5.2. In our experiments, we use a ring-shaped neodymium magnet (12 mm outside diameter, 5 mm inside diameter, 24 mm high). For any tool with this particular magnet, with a center-to-center distance between dipoles of 5 cm, we obtain a ratio of force per electrical current of 48 mN/A. This means the device can handle an averaged constant force of $F_r = 258$ mN ($P = 17$ W) with a peak force of up to $F_r = 959$ mN ($P = 230$ W) at full strength (using PWM control). This force value can be increased by increasing the volume of the tool magnet, with the trade-off of losing angular resolution and adding weight to the tool.




FIGURE 5.7: Electromagnet induced magnetization in each axis, $\mathbf{m}_e = (m_{e-x}, m_{e-y}, m_{e-z})$, as a function of the applied current settings (I_x, I_y, I_z). The magnetic field values are measured with hall sensors placed co-linear with each coil, and then transformed into M values.

5.4 USER EVALUATION

To assess the efficacy of our proposed approach, we validate the prototype in a perceptual study with 6 participants in order to 1) determine how well users can differentiate between different set-points, and 2) how accurate and precise users are with finding a set-point.

Procedure: Based on a pilot study, we predetermined 25 evenly separated set-points (Figure 5.8 right). We randomly selected a set-point, asked the user to find it, and report the corresponding number. Upon reporting, we also measured the Euclidean distance to the actual set-point. Every set-point was prompted exactly twice, resulting in 50 data points per user (300 in total). Only repulsive forces were tested. We used the same mapping parameters as in Figure 5.2.

Location accuracy: Figure 5.8 depicts the resulting confusion matrix between set-points. It can be seen that users accurately perceive discrete actuation points. For those actuation points that do cause incorrect answers, users tend to pick the neighboring location (typically higher on the sphere). This effect is pronounced along the meridian arc facing away from the user, whereas the orthogonal meridian produces fewer erroneous detections. This could be due to the position of the hand and arm and differences in muscle groups that are involved in actuating the wrist versus the whole hand. The difference in coil diameters could be another contributing factor.

Precision: We report the precision with respect to the angle θ . Figure 5.9 shows that the error increases as a function of the angle. A potential contributing factor here is




FIGURE 5.8: *Left*: confusion matrix of the 25 set-points, averaged over all users. High values on the diagonal indicate little confusion and the ability to differentiate between different set-points. *Right*: Set-points used in the study. The opacity directly correlates with the percentage of correct identifications by the users. Arrows are drawn when 33% or more of the *wrong* answers were attributed to set-point that the arrow points to.

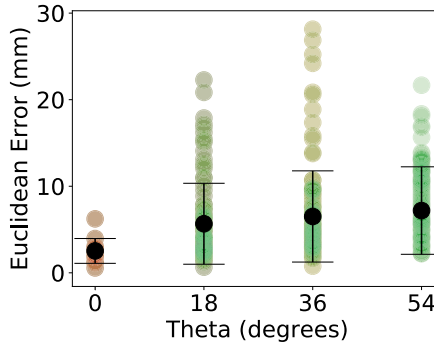


FIGURE 5.9: Euclidean distance between the true set-point position and the user reported position as a function of the azimuth (θ), measured from the top of the sphere and averaged over all angles and users.

that gravity has more impact on the pen the further down it moves on the hemisphere. This may make it more difficult for users to differentiate between the EM-actuation force and gravity. The mean errors of $2.5\text{mm} \pm 1.4$, $5.7\text{mm} \pm 4.6$, $6.5\text{mm} \pm 5.2$, and $7.2\text{mm} \pm 5.1$ are relatively small across the device.

5.5 DISCUSSION

In this chapter, we presented a novel contact-free volumetric haptic feedback device. This device uses a symmetric electromagnet combined with a dipole magnet model and a simple control law to deliver dynamically adjustable forces onto a handheld tool, such as a stylus. The tool only requires an embedded permanent magnet, allowing it to be completely untethered. Despite being contact-free, the force is grounded via the electromagnet, enabling the user to feel relatively large forces.

While our proposed method offers many advantages, it also has some drawbacks. Heat generation limits the number of interactions possible within a certain time frame. Specifically, when operating at full power, continuous interaction is limited to 5 seconds.

It is also important to note that the interaction between magnets involves both forces and torques. In this work, we focused on controlling the three force components via the 3 degrees of freedom (DoFs) of the electromagnet, allowing the torque values to adapt accordingly. However, the same procedure can be applied to control a specific torque map, leaving the force values unconstrained, or to manage a combination of force and torque. In future work, we aim to explore the dynamic capabilities of our approach, including advanced control schemes to continuously shape the force map.

Finally, our current control method relies on knowing the location of the tool, achieved through external cameras for optical tracking. However, this setup is cumbersome, expensive, and requires line-of-sight. In the next chapter, we will address this limitation by tracking the permanent magnet embedded in the tool using Hall sensors and a gradient-based method.

VOLUMETRIC SENSING AND ACTUATION OF PASSIVE MAGNETIC TOOLS FOR DYNAMIC HAPTIC FEEDBACK

In this chapter, we present Omni-v2, a self-contained 3D haptic feedback system capable of sensing and actuating an untethered, passive tool containing only a small embedded permanent magnet. Omni-v2 overcomes the limitations of Omni-v1 in two significant ways. First, it eliminates the need for external tracking by utilizing a novel gradient-based method to reconstruct the 3D position of the permanent magnet in midair. This is achieved using measurements from eight off-the-shelf Hall sensors integrated into the base. Second, Omni-v2 features an improved 3 DoF spherical electromagnet capable of delivering increased forces due to intertwined coils. The fully integrated Omni-v2 system, with no moving parts and no need for external tracking, is easy and affordable to fabricate. We detail Omni-v2's hardware implementation, our 3D reconstruction algorithm, and provide an in-depth evaluation of its tracking and actuation performance. Finally, we demonstrate its capabilities through a set of interactive usage scenarios.




FIGURE 6.1: We present Omni, a device that can simultaneously actuate and sense the position of a passive handheld tool. This is enabled through integrated hall effect sensors and our novel gradient-based optimization scheme. Omni can for example be used in 3D applications such as MR sculpting.

6.1 INTRODUCTION

This chapter builds on Contact-free Non-Planar Haptics with a Spherical Electromagnet (Chapter 5), by presenting *Omni-v2*, a self-contained 3D haptic feedback system. We enhance *Omni-v1* in two significant ways: first, by improving the actuator to enable increased peak and continuous forces, and second, by integrating sensing to overcome the limitations of expensive and cumbersome external tracking.

One of the core challenges in enabling high-fidelity free-space haptic interactions is to reliably track the input device in space while *simultaneously* exerting forces. While optical tracking systems can be used for 3D localization, they require significant instrumentation of the user’s environment, making them impractical for real-world deployments.

To address this challenge, we propose *Omni-v2*, an integrated system that can locate the tool in free space and deliver finely controllable haptic feedback via the same modality. Our system consists of an omnidirectional spherical electromagnet that delivers attractive and repulsive forces in radial and tangential directions onto a small, handheld magnet (Figure 6.1).

To track the tool’s location, we leverage eight integrated Hall sensors positioned in two separate XY-planes relative to the electromagnet’s coils. These sensor readings allow us to reconstruct the tool’s position based on the distortion of the magnetic field caused by the tool’s motion. Unlike previous approaches that rely on passive magnets [10, 155], our setting is more challenging due to the use of an active electromagnet. The electromagnetic actuation and sensing interfere, creating a superimposed magnetic field.

To decompose this superposition of magnetic fields, we introduce a novel 3D positioning algorithm. Our formulation reconstructs the magnet’s 3D position by minimizing the residual between the predicted magnetic field strength at a given location and the actual measurements from the Hall sensors. This approach is amenable to gradient-based optimization, and we demonstrate that the problem can be solved in real-time using a quasi-Newton solver, achieving an accuracy of 6.9mm after a one-shot calibration procedure.

We demonstrate the capabilities of our approach through a series of interactive use cases, some shown in Figure 6.10, leveraging the ability to track the handheld tool in 3D while simultaneously delivering dynamically adjustable haptic feedback. The magnet fits inside passive tools, such as a 3D-printed stylus, supporting fine-grained interaction tasks and alternative uses as a joystick. We showcase applications in gaming, accessibility, mixed reality, and 3D CAD design that utilize *Omni-v2*’s 3D tracking and actuation.

We detail our software and hardware implementation and thoroughly assess *Omni-v2*’s sensing and actuation capabilities. Our findings indicate that *Omni-v2* can continuously deliver up to 0.6N in any direction near the sphere, with the electromagnetic coils heating up to a sustainable 47°C. The system’s peak force reaches 1.8N for repelling and -3.1N for attractive forces to the sphere, and ± 2 N in the tangential plane. *Omni-v2* can actuate the magnet at 100Hz and sense and estimate its 3D position at 40Hz.

In summary, this chapter makes the following contributions:

- A fully self-contained system combining actuation and tracking components,
- A novel algorithm for the reconstruction of the 3D tool position under electromagnetic actuation,
- A thorough technical evaluation of the algorithm, achieving an accuracy of 5mm without actuation and 7mm with the electromagnet on,
- An assessment of the force generation capabilities, demonstrating up to 2N of haptic force in any direction,
- An exploration of use cases enabled by *Omni-v2*.

Our method can be extended to any apparatus involving a permanent magnet, integrating spatial position reconstruction and spatial actuation into a single device.

6.2 SYSTEM OVERVIEW

Omni-v2 is a self-contained haptic feedback system that simultaneously integrates 3D tracking and actuation using the same modality (see Figure 8.2). Through




FIGURE 6.2: Overview of our system. A 3D printed base contains the 3 DoF intertwined coils and the circular PCB with an array of eight hall sensors. Arbitrarily shapes tools can be 3D printed and augmented with a permanent magnet, to interact with *Omni-v2*.

actuating an untethered, contact-free tool by means of a magnetic field, our device supports rendering precise haptic attractive and repulsive forces as well as accurate tracking without the need for any external infrastructure or markers. Our system allows for rich interactions with and haptic perception of dynamic virtual surfaces.

Our goal is to enrich AR/VR and other 3D applications via *Omni-v2* and a minimally instrumented tool. Simplicity of the haptic prop and a walk-up-and-use experience were important design goals of our work. Furthermore, to create rich immersive experiences, such a system must be able to deliver different types of high-fidelity haptic forces and precisely sense user input without requiring any external tracking. Moreover, we aim for a self-contained device that is affordable and easy to manufacture.

The actuation mechanism used in *Omni-v2* is based on the working principle proposed in [286]. We build up a hemispherical shell base whose core houses a symmetric omnidirectional electromagnetic actuator. Three interwoven and mutually orthogonal coils generate the haptic forces.

By controlling the current in each coil, we can precisely configure the exerted force onto an external magnet, such as the one inside the 3D stylus tool. As the tool approaches the sphere, *Omni-v2* is able to provide independently controlled radial and tangential forces. Our design contributes several important improvements over [286] that allow us to provide twice the amount of force (2 N in either direction) and for much longer periods without suffering from self (over-)heating.

Although the tool is contact-free, the haptic force that the user perceives has its reaction force on the support base. In this sense, the user perceives grounded forces even if there is no mechanical link to the base.

In our current implementation, *Omni-v2* rests on a surface (e. g. table), though it is compact and could be mounted on a robotic end effector to deliver large-scale 3D feedback. This would allow for haptic feedback in a large volume, which would be

beneficial for VR applications. In this case, however, geomagnetism should be taken into account more strongly.

Beyond the improved actuator, our main contribution lies in the integration of the *Omni-v2* actuator with a fully self-contained real-time tracking method of the tool. To this end, eight Hall sensors are distributed below the interactive sphere. Each sensor reads a combination of the magnetic field generated by the tool, superimposed by the electromagnetic field generated by the actuator. We use a gradient-based optimization method to locate the magnet’s 3D position based on the Hall sensors’ readings, running at an interactive rate of 40 Hz.

6.3 METHOD

To support interactive experiences, our actuator needs to dynamically adjust its output according to the desired haptic feedback at a given time and tool position. We now describe our real-time approach to reconstruct the tool’s 3D position given the readings observed by the Hall sensors and the control strategy to govern the electromagnet-tool interaction.

6.3.1 3D position estimation

At the core of creating dynamic interactive experiences lies the ability to react in real-time to the movement of the user. Thus, a method to acquire the tool position with sufficient accuracy and precision with low latency is required. This is challenging due to the dynamic superimposition of the various magnetic fields (see Eq. 4.6). More precisely, directly computing the tool’s position from the sensor readings would require inversion of Eq. 4.1, which is non-linear and hence non-invertible, rendering an analytical solution for the tool position infeasible.

To overcome this difficulty, we introduce a reconstruction algorithm that optimizes an estimate of the tool’s position given the Hall sensor readings in real time. We propose an iterative model fitting approach for 3D position estimation. We minimize the residual between the *expected* sensor reading \mathbf{B}_i , as predicted by our model of the magnetic field (Eq. 4.6) given the current actuation, and the *actual* measurements acquired by the Hall sensors $\tilde{\mathbf{B}}_i$. In this setting, the optimization variables are the tool’s 3D position and its orientation. With a good initialization, which we attain by careful construction of the hardware, and exploiting the redundancy in the measurements, this algorithm provides accurate estimates of the tool’s position in 3D, with a mean accuracy of 6.9 ± 3.2 mm, as shown in our technical evaluation.

For each Hall sensor ($\mathbf{s}_i \in \mathbf{S}$) defined by its 3D coordinate, $\mathbf{s}_i = [s_x, s_y, s_z]^T$, we seek to find the tool position $\mathbf{r}_p = [p_x, p_y, p_z]^T$ and orientation $\mathbf{o} = [\theta, \varphi]^T$ that

provides the best model fit to the current reading. We use the global coordinate system for the sensors position \mathbf{r}_{s_i} and tool position \mathbf{r}_p with an origin in the center of the electromagnet, i. e. $\mathbf{r}_e = [0, 0, 0]$. The optimization problem is then given by:

$$\arg \min_{\mathbf{r}_p, \mathbf{o}} \left[\sum_{s_i \in S} \sum_{x,y,z} \mathbf{w}_i \left(\mathbf{B}_p(\mathbf{r}_{s_i} - \mathbf{r}_p, \mathbf{m}_p) + \mathbf{B}_e(\mathbf{r}_{s_i}, \mathbf{m}_e) + \mathbf{B}_n - \tilde{\mathbf{B}}_i \right)^2 \right], \quad (6.1)$$

where \mathbf{w}_i selectively weighs the sensor axes depending on the value it reads (i. e. a completely saturated sensor receives a weight close to 0). We pre-compute \mathbf{B}_e for different actuation strengths and \mathbf{B}_n denotes the background noise measured at startup. We empirically found that including the tool orientation θ and φ as free variables improves the position estimates by roughly 2 mm in Euclidean distance. However, the orientation estimates were too noisy to use in interactive settings.

We minimize Eq. 6.1 via iterative optimization. Specifically, we use PyTorch’s second-order L-BFGS optimizer, which typically works well for non-smooth optimization instances such as ours and requires no parameter tuning. Gradients are computed automatically via auto-grad.

Our method relies on known sensor locations obtained via one-shot calibration. We empirically found that an initial estimate of the sensors’ locations in the range of 1 mm accuracy is required to support robust convergence of the algorithm.

6.3.2 Actuation

Given the 3D pen-position, we can now deliver dynamically adjustable attractive and repulsive forces via the electromagnet to create desired haptic experiences. It remains in the hand of an application designer to decide with which intensity and in what direction the tool is pulled or pushed according to the desired user experience. *Omni-v2* is able to control 3 of the 6 DoF available, summing up forces and torques. We derive the case in which the goal is to control the three components of the haptic force \mathbf{F}_h , while controlling torques would follow an analogous derivation.

Under the magnetic dipole-dipole approximation, the force applied to the permanent magnet in the tool can be computed from the previous magnetic moments for the tool, \mathbf{m}_p , and the electromagnet magnetization \mathbf{m}_e we seek to control. Using the formulation of Yung et al. [285] and rewriting it in matrix form allows us to derive a simple control law for the parameters of the electromagnet \mathbf{m}_e^{set} , given the location of the tool \mathbf{r}_p , its dipole orientation \mathbf{m}_p , and the desired haptic force \mathbf{F}_h :

$$\mathbf{m}_e^{set} = a_1 [\mathbb{D} + a_2 \mathbb{I}]^{-1} * \mathbf{F}_h \quad (6.2)$$

where $a_1 = \frac{4\pi r_p^5}{3\mu_0}$, $a_2 = \langle \mathbf{m}_p, \mathbf{r} \rangle$ can be computed from the tool position information (as described in the previous section). The matrix \mathbb{D} has elements

$$d_{i,j} = m_{p_i} r_{p_j} + m_{p_j} r_{p_i} - 5 a_2 (r_{p_i} / r_p^2), \quad (6.3)$$

where m_{p_i} and r_{p_i} denote the i -component of the dipole \mathbf{m}_p (Eq. 4.2) and position \mathbf{r}_p , respectively. \mathbb{I} is a diagonal identity matrix.

Finally, we use the calibration matrix from Eq. 4.5 to find the electrical current to be applied to the individual coils:

$$\mathbf{I}^{set} = \mathbf{C}^{-1} * \mathbf{m}_e^{set}. \quad (6.4)$$

By combining Eq. 6.4 and 6.2, the vector of desired haptic force \mathbf{F}_h can be mapped into three actuation currents \mathbf{I}^{set} . This 3D-forces-onto-3D-currents mapping can always be decomposed into tangential and radial forces, using the tool's local coordinate system.

To only consider attractive and repelling forces and ignore the tangential component, Eq. 6.4 can be further simplified to:

$$\mathbf{I}^{set} = \alpha \frac{2\pi r_p^3}{3\mu_0} \mathbf{r}_p, \quad (6.5)$$

where α is the intensity of the force and its sign denotes attraction or repulsion relative to the sphere. In this particular case, the electrical current vector, the direction of the tool, and the tool dipole can be assumed to be always collinear.

6.3.3 Control implementation

The tracking algorithm, electromagnet control, and user-facing components (AR applications) run on a commodity gaming PC (Intel Core i7-8086K with 6 cores at 4 GHz, 32 GB RAM, NVIDIA GeForce GTX 1080 Ti) on Windows 10. The system is implemented in Python 3.7 and uses PyTorch's L-BFGS solver. The optimization-based tracking algorithm runs at 40Hz at the highest precision (6.9 mm). The AR applications were implemented in Unity 2019, SteamVR, and the Varjo Unity Plugin v2.4.

6.3.4 Hardware integration

Our hardware design is driven by two main factors. First, through a finite element analysis (FEA), we determined the physical characteristics of our hardware, such as coil diameter, core size, as well as the parameters of the permanent magnet embedded




FIGURE 6.3: *Omni-v2* hardware overview. Annotated view of the most important components of the system. All components can be acquired commercially or are easy to produce in a standard FabLab. A top-down view of the physical device can be seen in Figure Figure 6.1.

in the hand-held tool. This reference design strikes a balance between compact form-factor and force-generation capabilities. Second, we use off-the-shelf components for sensors and the voltage controller, in-house wound coils and in-house milled printed circuit boards (PCBs) to precisely mount the sensor's boards.

Reproducing *Omni-v2* only requires readily available components and few specialized tools, if any. The main hardware components are illustrated in Figure 6.3.

Figure 6.4 plots the results from an FEA of the tool's magnet to identify the best configuration given our current spherical electromagnet design. We use a magnet with 15 mm diameter and 7 mm height (volume of 5 cm^3), which corresponds to a point in the dark red region of the plot, where the output vertical force at 5 A is maximal. Please note that these characteristics hold for magnets with similar volume: either wider and shorter cylinders or narrow and tall, providing ample room for the design of the handheld prop. Importantly, a bigger magnet would not necessarily perform better and may decrease performance (top right region of the




FIGURE 6.4: Finite element analysis of actuation force as a function of the permanent magnet’s dimensions. Our current design is based on a 15×7 mm magnet, which provides maximal force generation capabilities. However, the plot illustrates that there is a Pareto front of similarly well-performing shapes that could allow for different tool designs.

plot), since the weight of the magnet counteracts the vertical actuation in repel mode. The omnidirectional electromagnet is based on a 30 mm diameter soft iron core encapsulated in 3D printed guides, aimed to assist during the manual winding (see Fig. 8.2). In contrast to [286], we construct our coils layer by layer in an interwoven fashion.

We iteratively add layers of x-winding, y-winding, and z-winding, respectively, until we reach an outer diameter of 65 mm. We use round copper wire with an external diameter of 0.9 mm (19 AWG) to obtain roughly 150 turns per coil. We employed a total cable length of $L_x = 21.4$ m, $L_y = 22.7$ m and $L_z = 24$ m for the x-, y- and z-coil respectively. Since the x-coil is wound first, it has a smaller radius per layer, i. e. a shorter perimeter per turn. The measured resistance of the coils are $R_x = 0.643 \Omega$, $R_y = 0.676 \Omega$ and $R_z = 0.708 \Omega$. The system supports up to 15 A of actuation current, which translates to a power of 157 W. To help remove the Joule heating generated within the coils, we place a brushless DC fan under the sphere (*CUI Devices*, $0.524 \text{ m}^3/\text{min}$), and include air intakes on the side of *Omni*.

To enable accurate and reliable tracking, it is paramount that the electromagnet and the Hall sensors are mounted rigidly with respect to each other. To ensure this, we fabricated custom PCBs using a desktop PCB milling machine (*Bantam Tools*). The ring-shaped sensor PCB is located below the electromagnet, with two circular arrays that mount 4 Hall Sensors each (*LIS3MDL*, *Pololu*).

All sensors are precisely aligned with the coil planes, such that each sensor’s local coordinate system aligns with the global frame. The Hall sensors sample at up to 1kHz and are read out by a microcontroller (*Teensy 4.0*), that communicates with the host PC.

We implemented an open-loop strategy to control the generated force. The approach is based on an analytical relation between the force, the coil actuation, and the tool location (Eq. 6.2), and relies on a few-point calibration. For actuation, we use three H-bridges (*Pololu G2 18v17*) to control the current of each coil with a pulse width modulation (PWM) of the voltage. Given Ohm’s law, we can directly control the current via the voltage if the resistance of each coil is known. Since electromagnets suffer from drift due to self-heating (and thus resistance changes), our system includes a coil-resistance drift compensation implemented directly on the actuation microcontroller. Two current sensors (*INA260*, *Adafruit*) provide an independent measure of the voltage and current of each coil. A sliding window average of the measured current I_i and voltage V_i are used to stabilize force-generation.

6.4 EVALUATION

Omni’s capability of delivering convincing haptic sensations relies on the performance of two main components: tracking of the tool position and in-air actuation. We performed technical evaluations on both aspects.

In summary, *Omni* is able to reconstruct the position of the tool with an accuracy of $6.9(\pm 3.2)$ mm and can deliver peak forces of up to ± 2 N, and 0.615 N continuously. Besides this technical evaluation, we demonstrate *Omni*’s interactive capabilities in the application section. We refer readers to Zarate et al. [286] for a psychophysics evaluation of a comparable underlying actuation mechanism, showing that users can discriminate at least 25 discrete force locations.

6.4.1 Tracking evaluation

To evaluate *Omni*’s tracking accuracy, we compared our position estimates to those of a 10-camera Optitrack setup, capturing a tracking space of 1.2×0.8 m with submillimeter accuracy at 100 Hz. We configured *Omni* to run in precision mode at a frame rate of 40 Hz. For both tracking methods, we recorded the position and rotation angles. We evaluated the accuracy of *Omni* in two conditions: *no actuation* and *actuation*. In the *no actuation* condition, no current was sent to the coils. In the *actuation* condition, the coils were actuated using a sawtooth function sending current between -4 and 4 Amperes for each axis.

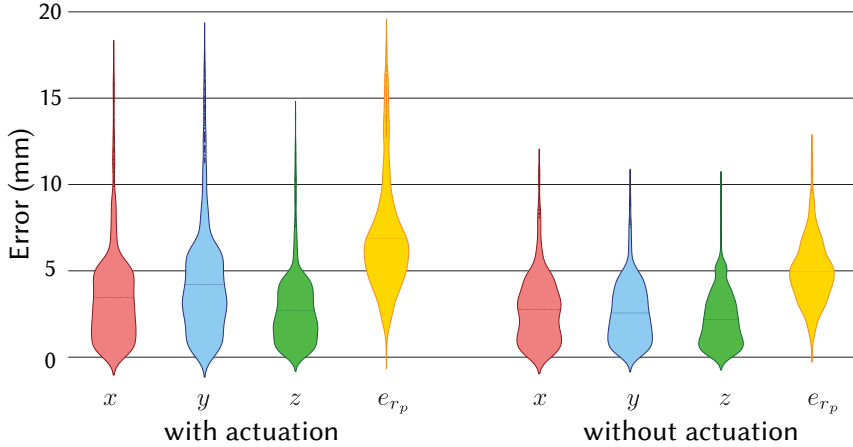


FIGURE 6.5: Distribution of tracking error with and without current applied to the electromagnet.

For each condition, the pen was moved around the center of *Omni* at a distance of up to 10 cm, covering the area around the device. We collected 1600 samples for the *no actuation* condition and 2600 samples for the *actuation* condition, both at roughly 5 Hz.

6.4.1.1 Results

We found that the average difference between the two tracking systems is $e_{r_p} = 4.9(\pm 1.8)$ mm in the *no actuation* condition and $e_{r_p} = 6.9(\pm 3.2)$ mm in the *actuation* condition. Analyzing each axis separately, we found that $\mathbf{e}_{r_p} = [3.4; 3.1; 2.7]$ mm for the tracking errors and *no actuation* and $\mathbf{e}_{r_p} = [4.4; 5.2; 3.3]$ mm in the *actuation* condition. The results are summarized in Figure 6.5.

Finally, we tested our formulation with and without the orientation estimation of the magnet. While we found that including these additional optimization variables improves the accuracy of the position estimates, these estimates are unstable and not yet useful for interactive applications.

Intuitively, it makes sense that including the orientation in the model fitting improves position estimation since the orientation of the magnet does influence the magnetic field. Furthermore, it is known that the model we leverage [285] works best for spherical magnets (e. g. point estimates of positions) and hence the models' approximation error may be a source of noise in the orientation estimates. We leave an extension of the reconstruction method to 5-DoF for future work.




FIGURE 6.6: Setup for actuator evaluation. An additional 3D printed hemisphere is placed on top of *Omni* to hold the force sensors.

6.4.2 Actuator evaluation

Omni's 3 DoF spherical electromagnet produces a force on the permanent magnet in the tool by dynamically adjusting the magnetic field through currents in the orthogonal coils. To quantify this actuation, we measured the radial and tangential forces at different locations around the electromagnet in *Omni*'s spherical base. We placed a 3D-printed hemisphere over the electromagnet (see Figure 6.6). The hemisphere has three slots to place a test magnet (*S-30-07-N*, *Supermagnete*, same as in tool) and two force sensors (*FSAGPNXX1.5LCAC5*, *Honeywell*). The force sensors were attached between the electromagnet and the test magnet to measure radial force, and to the side of the test magnet to measure tangential forces.

6.4.2.1 Results

We generally observed a linear response of actuation with respect to the applied current. On top of the electromagnet, we measured a maximum vertical repulsive (radial) force of 1.95 N at $I_z = 14.6$ A and a maximum attractive force of -3.04 N at $I_z = -14.6$ A, shown in Figure 6.7.

When *Omni* applies $I_z = +3.7$ A, it compensates for the weight and snapping and the magnet starts to levitate¹. Note that at this position, the force is the sum of the electromagnetic actuation, the snapping to the core, and the gravitational attraction.

¹ In this paper we use the term *levitation* in the sense of *compensate its weight completely*. A complete levitation would need to control the actuation in all three axes to keep the magnet floating in place.




FIGURE 6.7: Radial and tangential forces on the permanent magnet as a function of coil actuation I_x , I_y and I_z , for the magnet located on top of the sphere (Fig 6.6). Force was collected with a compression-like force sensor.

The weight of the tool produces a force of $F_r = -370$ mN (38 gr), while ferro-magnetic snapping yielded additional 170 mN of force, combined these produce an attracting radial force of $F_r = -540$ mN without actuation ($I_z = 0$ A). All those components contribute to users' perception of force.

On top of the sphere, the weight and snapping are orthogonal to the x -axis and y -axis and do not influence the radial forces along those axes. Consistently, we measured a linear response on those axes of the form $F_{r_{x-axis}} = 0.122 [N/A] I_x$ and $F_{r_{y-axis}} = 0.142 [N/A] I_y$. For the other locations in our test setup (horizontal to vertical), we observed forces in the range of ± 2 N at ± 15 A with the corresponding corrections for weight and snapping. Note that the forces have been measured when the tool was in contact with *Omni's* hemisphere.

The force intensity decays with $1/(d_0 + g)^4$, where g is the air-gap between the tool and the sphere. The parameter $d_0 = 41$ mm is the center-to-center distance between the electromagnet and the permanent magnet when the tool touches the sphere. For example, for $g = 10$ mm, the reachable range of forces drops to ± 835 mN. At 30mm from the hull this reduces to 0.2N

6.4.2.2 Evaluation of EM heating

To test the stability of the generated forces and the thermal capabilities of our system, we ran two experiments. First, we set the y -axis coils to maximum actuation current




FIGURE 6.8: Temporal evolution of the self-heating of the coils for two different types of actuation. Top: a *peak-force* of 2 N ($I_y = 15$ A) during 25 seconds. Bottom: a *constant-force* of 600 mN ($I_y = 5$ A) during 15 minutes.

$I_y = 15$ A for 25 seconds and let it cool down afterwards to test the system under *peak-force* conditions. Second, we set the same coil to 1/3 of the maximum actuation and we let it run for 15 minutes, to test under *constant-force* conditions. Figure 6.8 shows the evolution of the generated force and the temperature of the coil for both conditions.

During *peak-force*, the system delivers a force of 2.04 ± 0.04 N. Starting from room temperature (24 °C), the actuator heats up to 39 °C but only 40 seconds after the actuation has been turned off, showing the system's thermal inertia. The $\Delta T = 15$ °C during this intense actuation peak shows that our system is capable of thermally buffering and dissipating the heat generated by intense forces even during tens of seconds.

In our *constant-force* experiment (Figure 6.8, *bottom*), the force remained constant within the limits 0.615 ± 0.015 N and for a duration of 15 min, even when the temperature of the coils (and their resistances) significantly changed. In addition to compensating for the actuation drifts, we used the coils' resistance changes over time as the limiting factor to avoid overheating of the coils and the 3D printed parts, in case the system is required to apply maximum forces for minutes.



FIGURE 6.9: We combine *Omni* with a video see-through MR device (Varjo XR-1) to showcase the applications.

6.5 APPLICATIONS

We implemented a series of applications to showcase *Omni*'s potential in supporting spatial interaction with virtual objects that is supported by strong haptic sensations as shown in Figure 6.10. While *Omni* can support traditional desktop interaction with haptic cues (e. g. free-form tool-based gesture input), we focus our applications on mixed and virtual reality scenarios that are inherently spatial.

Specifically, we demonstrate scenarios in MR, which benefit most from *Omni*'s walk-up-and-use nature to track and haptically actuate an untethered, small magnet in the space around *Omni*'s base.

We implemented our applications using a video pass-through Mixed Reality device (Varjo XR-1), shown in Figure 6.9. *Videos and photos were recorded live through Varjo's software.*

6.5.1 Sculpting

Figure 6.10 illustrates how *Omni* haptically supports 3D sculpting and CAD design.

Here, the user finely selects locations on a 3D base object for extrusion by means of the stylus, which is tracked through *Omni*. When extruding individual bumps from the starting configuration, the user can probe and feel the compliance of the material, rendered through attractive and repulsive forces. Having extruded several




FIGURE 6.10: We present possible use cases for *Omni*. *Left* shows the possibilities in 3D CAD design, in this case sculpting. *Center* shows a user exploring and manipulating an augmented reality object. For both applications, users can feel the shape of the outer hull of the objects. *Right* shows a racing game. Once the car collides with the wall, the pen gets pushed to the base. *Arrows indicate movement and are drawn on top of the photo to increase clarity.*

bumps from the original shape, the user may inspect the 3D object visually as well as haptically, as *Omni* renders collisions with the tool through tangential actuation.

Following this 3D interaction scenario, *Omni*'s haptic capabilities could be scaled to common 3D editing techniques such as grid snapping, guided object rotation, and 3D transformation.

6.5.2 *Non-rigid object exploration*

Omni's tracking and actuation also lends itself to haptically rendering geometric objects that are non-rigid and may have anisotropic material properties, such as geographical surfaces, enlarged microscopic surfaces, or other complex geometries. We demonstrate how *Omni* generates haptic feedback while touching and poking a virtual dragon that is configured to simulate rubber-like material properties. Here, the force *Omni* renders increases with the amount of object deformation, which portrays the *physical behavior* more accurately than would be possible to experience through mere visual feedback.

6.5.3 *Gaming*

Finally, we demonstrate how *Omni* can be used for enhancing the experience of gameplay. Using the magnet-equipped tool as a joystick, we demonstrate how users can steer a car in an AR racing game. The combination of *Omni*'s haptic feedback and visual control over the car increases the level of immersion provided by the

system, as haptic and visual sensations render a coherent experience. For example, *Omni* renders car collisions with forces whose magnitude depends on collision speeds.

6.6 DISCUSSION

Omni-v2 is capable of tracking a passive tool with an accuracy of roughly 6.9 mm and, at the same time, deliver a maximum force of up to 2 N to the tool. This is enabled by our novel gradient-based approach in 3D position reconstruction that accounts for the force exerted by the electromagnet.

Over extended periods of time, *Omni-v2* can comfortably produce a force of 0.615 N without the risk of overheating. In our applications, we show that *Omni-v2* has the potential for a wide range of usage scenarios, specifically to enrich AR and VR interactions.

Omni-v2 is, however, not limited to spatial applications. We believe that *Omni-v2* can be a valuable addition to desktop interfaces, e.g., navigating through video editing tools or gaming. We plan to broaden *Omni-v2*'s usage scenarios in the future.

The overall tracking performance of *Omni-v2* suffices for interactive applications such as the ones shown in this paper. The accuracy could be improved by adding more Hall sensors, or optimizing their placement further (e.g., placing them on the outer hull of the device). Furthermore, a spherical tip on the passive tool that more closely resembles the dipole in our magnetic model could further improve *Omni-v2*'s accuracy. We believe, however, that the design of *Omni-v2* represents a good balance of cost and complexity of manufacturing, and accuracy.

Our current implementation of *Omni-v2* and the accompanying tracking and actuation algorithms assumes the presence of a single passive tool. Our method, however, potentially generalizes to tracking multiple passive tools by accounting for the presence of multiple permanent magnets. This poses another interesting challenge: the magnets of multiple tools will interact with each other, i.e., attract and repel each other. The electromagnet will also jointly interact with those tools, leading to challenges in terms of computation and convergence. We believe that our gradient-based optimization can account for such interactions and plan to investigate this in the future.

In developing and testing our applications, we found that *Omni-v2*'s current frame rate of 40 Hz suffices for many interactive scenarios. The frame rate is a trade-off between speed and accuracy. In our tests, decreasing the desired accuracy in our optimization doubled the frame rate, while resulting in errors in the 3D position estimation of more than 1 cm, however. Finding the sweet spot for this trade-off depends on the application. While our applications worked well with 40 Hz and

the current accuracy, more intricate actions such as high-precision sculpting might benefit from higher frame rates *and* precision. Reducing the latency of several system components (e.g., sensor latency, convergence time) is another interesting direction of future research.

Furthermore, the control strategy we used was fairly naïve, as it only takes the current tool position into account. A model predictive strategy could account for future states, user intent, and optimize to reduce heating. We will explore in the next chapter how model predictive approaches can be used for haptic systems.

Overall, the main benefits of *Omni-v2* lie in the high accuracy and large force it can produce. It does so without mechanically moving parts, which would be subject to wear. Such wear is not the case for our device, because it is exclusively based on electromagnetic force. We believe that different form factors of *Omni-v2* (e.g., body-mounted, larger size) can present interesting directions of future research. A body-mounted version could be interesting for VR applications in which the user moves in 3D space. The larger size could result in more discernible points.

Additionally, the influence of strength on user perception and factors such as just-noticeable-difference will allow us to characterize the benefits and challenges of *Omni-v2*, and electromagnetic haptic devices in general. We believe that *Omni-v2* opens interesting directions for future research in terms of novel devices, and magnetic actuation and tracking.

In this chapter, we introduced *Omni-v2*, a novel electromagnetic platform that simultaneously tracks and actuates a permanent magnet in the space around it. Our self-contained base assembly integrates 3D magnetic sensing using Hall sensors and magnetic actuation through radial and tangential forces produced by three orthogonal electromagnetic coils within a single sphere.

Our core contribution lies in decomposing the natural interference caused by simultaneous magnetic tracking and actuation. This is enabled by our novel gradient-based optimization method, which minimizes the difference between estimated and observed magnetic fields. This approach affords 3D tracking capabilities with a mean error of 6.9 mm during actuation forces of up to 2 N.

Omni-v2's capabilities allow spatial interaction systems to integrate 3D tracking and actuation of untethered, free-ranging tools simply by embedding a small permanent magnet. We demonstrated a series of example applications leveraging *Omni-v2*'s capabilities, showcasing its potential in various interactive scenarios.

However, an open question remains on how to control the system to enable user autonomy rather than guidance. Addressing this challenge will be a key focus for the next part of this dissertation, aiming to further enhance the usability and functionality of haptic devices in practical applications.

SUMMARY & INSIGHTS

SUMMARY We introduced two novel electromagnetic platforms designed to actuate a permanent magnet in the surrounding space. The first project, "Contact-free Non-Planar Haptics with a Spherical Electromagnet" (Chapter 5), focused on designing and fabricating a spherical electromagnet capable of generating adjustable radial and tangential forces using three orthogonal electromagnetic coils integrated into a single sphere. This system targets a handheld tool, such as a stylus, and allows users to feel relatively large forces despite the tool being untethered. A user experiment with six participants characterized the force delivery and perceived precision, revealing a resolution of at least 25 discernible locations for repelling forces distributed across the hemisphere. However, *Omni-v1* required external tracking of the tool, which was cumbersome and expensive.

In "Volumetric Sensing and Actuation of Passive Magnetic Tools for Dynamic Haptic Feedback" (Chapter 6), we addressed the tracking limitation of *Omni-v1* by integrating 3D magnetic sensing using Hall sensors with magnetic actuation through an improved spherical electromagnet. This integration allows spatial interaction systems to incorporate 3D tracking and actuation of untethered tools by embedding a small permanent magnet. The core contribution of *Omni-v2* is resolving the interference caused by simultaneous magnetic tracking and actuation using a novel gradient-based optimization method. This method provides 3D tracking capabilities with a mean error of 6.9 mm during actuation forces of up to 2N. We demonstrated several example applications showcasing *Omni-v2*'s capabilities.

IMPLICATIONS In our context, both the user and the system interact with a shared variable, such as car acceleration or a graphical menu. Kinesthetic haptic feedback is a special case where this variable (e.g., joystick position) is both acted upon and perceived through the same modality, like the user's hands. Previous research on kinesthetic haptic feedback generally focuses on complex mechanical devices (e.g., [9, 177, 249, 253, 267, 291]). Extensions of these systems include exoskeletons [44, 92], gloves [55, 99], and tilt-platforms [128, 217]. However, these often necessitate user instrumentation and introduce system friction that users can always feel, even when the system is off, thereby limiting full user control.

In contrast, our work provides grounded yet untethered forces onto a tool, thereby overcoming the limitations of not covering the full spectrum from user agency to system automation. Our work lays the foundation for investigating haptic devices

from a new perspective, that of shared variable interfaces. We have taken the first step towards understanding these interfaces, demonstrating that haptic shared variable interfaces open up novel applications in gaming, object exploration, and design. By eliciting forces on the tool manipulated by the user, we enable haptic feedback in the same modality as the user's exertion, resulting in natural interactions. This integration of sensing and actuating the tool through electromagnetism is both cost-effective and efficient.

LIMITATIONS Both *Omni-v1* and *Omni-v2* have limitations within the context of interfaces with shared variables.

First, while our work suggests that interacting with a shared variable haptic interface is natural for users, it remains uncertain to what extent our findings generalize to other interface types, such as graphical interfaces. Further investigation is needed to understand the interaction implications beyond the haptic paradigm.

Second, we have not compared our results to haptic feedback without a shared variable, such as vibrotactile feedback. Comparing against a broader range of haptic devices would allow us to draw more comprehensive conclusions regarding usability improvements.

Finally, we have not explored an optimal control strategy for haptic feedback. Haptic feedback involves not only the device itself but also how the device is controlled. Currently, we provide haptic feedback based on heuristics and the current tool position. An adaptive strategy incorporating a predictive user model might be more effective. Such a control method could balance autonomy and automation, and account for user intent, making the device more dynamic.

Part III

THE CONTROL OF SHARED VARIABLES
INTERFACES

OPTIMAL CONTROL FOR ELECTROMAGNETIC HAPTIC GUIDANCE SYSTEMS

In Chapter 5 and Chapter 6, we introduced novel actuator and sensing techniques for electromagnetic haptic devices. However, the challenge of effectively controlling the actuator given the sensor input remains unresolved. In this chapter, we address this by exploring an optimal control method for electromagnetic haptic guidance systems. Our approach assists users in pen-based tasks such as drawing, sketching, and designing, while ensuring that user agency is maintained. Traditional methods force the stylus to follow a continuously advancing setpoint on a target trajectory, often resulting in loss of haptic guidance or unintended snapping. In contrast, our control approach gently pulls users towards the target trajectory, allowing for spontaneous adaptation and drawing at their own speed. To achieve this flexible guidance, we iteratively predict the motion of an input device (such as a pen) and dynamically adjust the position and strength of an underlying electromagnetic actuator. To enable real-time computation, we introduce a novel, fast approximate model of an electromagnet. We validate our approach on a prototype hardware platform featuring an electromagnet on a bi-axial linear stage and demonstrate its effectiveness through various applications. Experimental results indicate that our method is more accurate and preferred by users compared to open-loop and time-dependent closed-loop approaches.




FIGURE 8.1: We propose an optimal control scheme for electromagnetic guidance systems (left). A target trajectory is provided, on which users are guided. They can always adapt the trajectory, our optimization then guides users back to the target (offset between target and drawing for illustration purposes).

8.1 INTRODUCTION

In this dissertation, we have presented pen-based haptic control systems, including a novel spherical electromagnet (Chapter 5) and a sensing algorithm (Chapter 6). However, while these systems aim to control the pen, they inadvertently reduce user agency. A haptic system should guide the user while preserving their autonomy.

Pen-based interactive systems featuring *haptic guidance* support users in various applications such as drawing, sketching, writing, and CAD design. The goal of such systems is to enable users to draw higher-complexity shapes with less effort and higher accuracy, or to support users through virtual haptic tools such as rulers or guides. Crucially, such systems aim to strike a balance between giving users a strong sense of control and agency, while providing feedback unobtrusively. This is enabled through embedding controllable electromagnets in the system, which can then guide input devices such as a pen, or provide feedback to users about the positions and boundaries of virtual objects.

Existing systems such as dePenD [278] typically employ an *open-loop control approach*. The magnet that drives the pen is set to a predefined trajectory, and users then must closely follow the movement of the system. In this case, it is not possible for users to adjust the trajectory, since it would lead to a loss of haptic guidance. This effectively leads to a decrease in control of users, and arguably a loss of user agency.

Alternatively, it is possible to extend haptic guidance systems with traditional closed-loop approaches, e. g. by implementing a proportional–integral–derivative (PID) controller [12] (and as we did in *Omni-v1* and *Omni-v2*) or based on heuristics [165]. Such systems adjust to users' movement but are, usually, based on a timed reference, effectively dictating users' drawing speed. This can lead to unintended behavior such as snapping whenever the pen is too close to the actuator, a problem

that is exacerbated for magnetic systems due to the non-linear nature of the magnetic force over distance.

We propose a real-time closed-loop control approach that allows users to retain agency and control while being assisted by an electromagnetic haptic guidance system. Our approach enables users to draw at their desired speed and adjust their target trajectory continuously. It then adapts and complies with such modifications while giving corrective feedback. Our algorithm positions and regulates a variable-strength electromagnet such that it provides dynamically adjustable in-plane magnetic forces to the pen tip.

We contribute a novel optimization scheme for electromagnetic-based haptic guidance systems, i. e. models and control algorithm, that enables formalizing this problem in the established Model Predictive Contour Control (MPCC) framework [140], which has previously only been employed in contexts such as RC-racing [159] or drone cinematography [193].

We provide an accurate system model, parameters, and an appropriate cost function alongside a method to optimize the model parameters given user inputs. Modeling the non-linear interaction of an electromagnetic force field typically makes use of the finite element method (FEM), which is not applicable for real-time scenarios. To overcome this challenge, we additionally contribute a novel approximate yet accurate model of the electromagnetic force field that can be evaluated analytically in real time. Compared to simpler control schemes such as Model Predictive Control (MPC) [70] and many implementations of PID control, our approach does not require a timed reference and hence allows users to draw at their desired speed.

Furthermore, our optimization scheme allows for error-correcting force feedback, gently pulling the user back to the desired trajectory rather than pushing or pulling the pen to a continuously advancing setpoint on the trajectory. With our approach, the reference path can be updated at every timestep, thus allowing users to continuously change their desired trajectories. This enables applying the algorithm to fully dynamic references, for example virtual tools such as rulers or programmable French curves.

To assess the proposed control algorithm, we developed a proof-of-concept hardware implementation (see Figure 8.2), leveraging an electromagnet that moves underneath the drawing surface or display on a bi-axial linear stage. The magnet provides variable strength guidance onto the tip of a minimally instrumented pen or stylus via an electromagnet positioned directly below a drawing surface, guided by our proposed approach.

We demonstrate the feasibility of our approach with a set of applications, specifically drawing guidance on conventional paper for sketching and writing, and a digital sketching application that features virtual haptic guides and rulers.




FIGURE 8.2: We implement our proposed guidance system using a two-axis linear stage equipped with an electromagnet. All technical and user evaluations were completed using the Pressure Sensitive Tablet (*right*). We additionally developed an all-digital implementation using a multi-touch tablet with display (*left*).

To evaluate our approach, we performed two experiments with twelve participants each. We first compared free-hand drawing of shapes with varying complexity with and without our feedback system. Results showed that the haptic guidance using our approach improved the accuracy across shapes by up to 50% to 1.87 mm. We then compared our approach to our implementation of dePENd (open-loop) and a simple MPC-based closed-loop control scheme. Our approach showed significantly higher accuracy and was preferred by users.

In summary, we contribute

- A novel MPCC-based optimization scheme for electromagnetic haptic guidance systems including models, parameters, cost function and control algorithm.
- A novel real-time approximate model for electromagnets that generalizes beyond our hardware implementation.
- Evaluations showing the improved accuracy of our method.

8.2 METHOD OVERVIEW

The goal of our online optimal control scheme is to allow users to maintain control and agency over the input device (e.g., pen, stylus), while experiencing dynamic guidance from the system. Importantly, it leverages *time-free references*, and thus the dynamics are entirely driven by the pen position over time, which is different from approaches such as MPC.

The proposed optimization scheme allows us to adjust the magnet position and strength such that it gently pulls the pen tip towards a desired stroke, while allowing users to draw at their desired speed and without fully taking over control. The algorithm is generally hardware agnostic and works for devices with electromagnetic actuators underneath an interaction surface. This can be implemented via bi-axial linear stage as in our prototype (see Figure 8.2) or via a matrix of electromagnets which would lend itself better to miniaturization. Furthermore, the algorithm requires a reference trajectory over the optimization horizon. This can be defined a priori, such as a known shape to be traced, or may be provided dynamically, e.g., the output of a predictive model (e.g., Aksan et al. [7]).

At each time step, we minimize a cost functional over a receding time horizon to find optimized values for system states \mathbf{x} and inputs \mathbf{u} . As a high-level abstraction, the cost function

$$\underset{\mathbf{x}, \mathbf{u}}{\text{minimize}} \underbrace{\sum \mathcal{C}_{\text{force}}(\mathbf{x}, \mathbf{u})}_{\text{Eq. 8.3, 8.4 \& ??}} + \underbrace{\mathcal{C}_{\text{path}}(\mathbf{x}, \mathbf{u})}_{\text{Eq. ??}} + \underbrace{\mathcal{C}_{\text{progress}}(\mathbf{x}, \mathbf{u})}_{\text{Eq. ??}}, \quad (8.1)$$

serves three main purposes: 1) ensuring that the user perceives haptic feedback of dynamically adjustable force ($\mathcal{C}_{\text{force}}$), 2) stays close to the desired path but does not rigidly prescribe it ($\mathcal{C}_{\text{path}}$), and 3) makes progress along it ($\mathcal{C}_{\text{progress}}$) but allows the user to vary drawing speed freely.

8.3 METHOD

Our main contributions are models and a control strategy that enables using the MPCC framework [141] for electromagnetic haptic guidance. MPCC is a closed-loop *time-independent* control strategy that minimizes a cost function over a fixed receding horizon. There are several advantages in using our formulation over open-loop (as used in dePENd [278]) or time-dependent strategies (e.g. MPC). First, closed-loop control allows to react to user-input, whereas open-loop control removes all user agency. Both MPC and MPCC are closed-loop control strategies. However, MPC tracks a timed reference, requiring a fixed velocity by users. MPCC follows a time-free trajectory, which allows the user to progress at their own speed. Figure 8.3

Name	Range / Value	Description
\mathbf{p}_p	\mathbb{R}^2	Position of pen
\mathbf{p}_e	\mathbb{R}^2	Position of electromagnet
\mathbf{F}_a	\mathbb{R}^3	Electromagnetic force vector
α	$[0, 1]$	Electromagnetic intensity
\mathbf{s}	$\theta \in [0, L]$	Target trajectory of length L
\mathbf{x}	$[\mathbf{p}_m, \dot{\mathbf{p}}_m, \alpha, \theta]$	System states
\mathbf{u}	$[\dot{\mathbf{p}}_m, \dot{\alpha}, \dot{\theta}]$	System inputs

TABLE 8.1: Overview control parameters and values




FIGURE 8.3: Overview of different control strategies on a target trajectory (green), with constant pen position. For open-loop, the position of the electromagnet is identical to the constantly advancing setpoint, leading to loss of haptic guidance. For MPC, although the pen is static, the guidance changes at every timestep since the setpoint advances. In our approach, the setpoint is also based on the pen position, therefore remains stationary in this case and guides the user towards the target trajectory.

illustrates the expected behavior for the different strategies, given that the user slows down or stops moving the pen. The desired behavior here would be that the algorithm essentially “waits”, i. e. provides guidances towards a slowly or no-longer advancing setpoint. In this situation, open-loop approaches would lead to lost haptic guidance. Closed-loop time-dependent approaches would guide the pen towards a constantly advancing setpoint (although users do no longer move), which can lead to problems such as the user being guided backwards (e. g. timestep $t = 3$ is in front of $t = 2$).

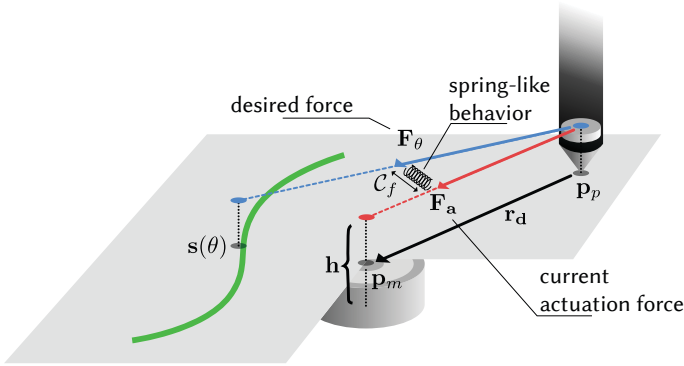


FIGURE 8.4: Illustration of actuation force \mathbf{F}_a , desired force \mathbf{F}_θ , and the force cost-term \mathcal{C}_f associate with the difference between those two forces.

Our method is designed to exert a force \mathbf{F}_θ of desired strength onto the pen to guide the user towards the target trajectory \mathbf{s} . The path \mathbf{s} of length L is parametrized by $\theta \in [0, L]$. Note that we do not prescribe how fast users draw and hence for each given pen position \mathbf{p}_p we first need to establish the closest position on the path parameterized by $\mathbf{s}(\theta)$. The vector between the pen position and $\mathbf{s}(\theta)$ is defined as \mathbf{r} . We leverage a receding horizon optimization strategy and the global reference can hence be adjusted or replaced entirely at every iteration. The path \mathbf{s} is then a local fit to the global reference. Furthermore, we seek to find optimized values for the electromagnet intensity α and the in-plane electromagnet position \mathbf{p}_m . Solving the error functional given in Eq. (8.10) at each timestep yields optimized values for system states \mathbf{x} and inputs \mathbf{u} .

As common in MPC(C), the system is initialized from measurements at $t = 0$. The system state is then propagated over the horizon with the dynamics model $f(\mathbf{x}, \mathbf{u})$. The system state vector \mathbf{x} contains variables that are controlled by the algorithm (magnet intensity and position, current path progress). The first of the optimized inputs (u_0) is then applied to the physical system, transitioning the system state to x_1 , before iteratively repeating the process to allow for correcting modeling errors.

8.3.1 Haptics model: controlling the force of the electromagnet

The main goals of our approach is that users can move freely in terms of position and speed, and that the actuator continuously pulls them towards an advancing setpoint $\mathbf{s}(\theta)$ on the target trajectory \mathbf{s} . At any time, the magnet exhibits an actuation force \mathbf{F}_a on the pen, given by our electromagnetic force model (see Implementation section).

Therein lies the challenge, illustrated in Figure 8.4. The setpoint is continuously advancing based on the movement of the pen to ensure progress. The actuator needs to pull the pen towards the setpoint by exhibiting force \mathbf{F}^* , but currently exhibits \mathbf{F}_a . The two forces only align if the pen is exactly at the setpoint, which is rarely the case. To overcome this challenge, we propose modeling this interaction by a spring-like behavior that “pulls” \mathbf{F}_a towards \mathbf{F}^* . In this way, the magnet continuously guides the pen towards the setpoint, and the force linearly increases with distance between the pen and the target setpoint denoted as:

$$\mathbf{F}_\theta(\mathbf{r}) = c F_0 \mathbf{r} \cdot \mathbf{e}_r. \quad (8.2)$$

Here \mathbf{e}_r is a unit vector in the direction of \mathbf{r} , c is a scalar that regulates the stiffness of the spring (in our case $c = 5/h$), F_0 a scaling of the EM force (i. e. the force felt by users) and h the distance between dipoles in z (see Fig. 8.4). Although simple, this formulation ensures that the haptic guidance is strong under large deviation from the path while vanishing as the user approaches the target path ($r_\theta \rightarrow 0$). Note that Eq. 8.2 is a design choice. Different formulations can be used to achieve different user experiences. Furthermore, replacing our hardware prototype and force-model would allow for adaptation of the remainder of the method to different actuation principles.

The above haptics model serves as basis for our problem formulation of electromagnetic guidance in the MPC framework. Using the vectors of the current actuation force \mathbf{F}_a and desired force \mathbf{F}^* , we formulate a quadratic cost term to penalize the difference between desired force and actual force as:

$$\mathcal{C}_f(\mathbf{p}_e, \mathbf{p}_p, \alpha) = \|\mathbf{F}_\theta(\mathbf{r}) - \mathbf{F}_a(\mathbf{d})\|^2. \quad (8.3)$$

where \mathbf{d} is the in-plane vector between the magnet and the pen. Since the actuation force \mathbf{F}_a declines rapidly with distance \mathbf{d} , the gradient of \mathcal{C}_f goes to 0 for large values of \mathbf{d} causing the optimization to become unstable. To counterbalance this issue we encourage the electromagnet to stay close to the pen:

$$\mathcal{C}_d(\mathbf{p}_e, \mathbf{p}_p) = d^2. \quad (8.4)$$

Finally, we prioritize proximity between the magnet and the pen rather than increasing its force by penalizing excessive use of magnetic intensity α :

$$\mathcal{C}_\alpha(\alpha) = \alpha^2. \quad (8.5)$$

8.3.2 Controlling the position of the electromagnet

We continuously optimize the position of the electromagnet with the goal of keeping the distance between the desired path and the pen minimal. To give the user freedom




FIGURE 8.5: Illustration of lag- and contouring error decomposition.

in deciding their drawing speed we first need to find the reference point $\mathbf{s}(\theta)$ on the target trajectory \mathbf{s} . Finding the closest point on the path is an optimization problem itself and hence can not be used within our optimization. Similar to recent work in robot trajectory generation [84, 193], we decompose the distance to the closest point into a contouring and lag error, as shown in Figure 8.5. \mathbf{r} is the vector between the pen \mathbf{p}_p and a point $\mathbf{s}(\theta)$ on the spline, and \mathbf{n} as the normalized tangent vector to the spline at that point, which is defined as $\mathbf{n} = \frac{\partial \mathbf{s}(\theta)}{\partial \theta}$. The vector \mathbf{r} can now be decomposed into a lag error and a contour error (Figure 8.5). The lag-error \mathcal{C}_l is computed as the projection of \mathbf{r} . The contour-error \mathcal{C}_c is the component of \mathbf{r} orthogonal to the normal:

$$\begin{aligned}\mathcal{C}_l(\mathbf{p}_p, \theta) &= \|\langle \mathbf{r}, \mathbf{n} \rangle\|^2, \\ \mathcal{C}_c(\mathbf{p}_p, \theta) &= \|\mathbf{r} - (\langle \mathbf{r}, \mathbf{n} \rangle) \mathbf{n}\|^2.\end{aligned}\tag{8.6}$$

Separating lag from contouring error allows us, for example, to differentiate how we penalize a deviation from the path (\mathcal{C}_c), versus encouraging the user to progress (\mathcal{C}_l). We furthermore include cost terms to ensure that the magnet stays ahead of the pen (\mathcal{C}_θ) and to encourage smooth progress ($\mathcal{C}_{\dot{\theta}}$) computed as

$$\begin{aligned}\mathcal{C}_\theta(\theta) &= -\theta, \\ \mathcal{C}_{\dot{\theta}}(\dot{\theta}) &= (\dot{\theta}_t - \dot{\theta}_{t-1})^2.\end{aligned}\tag{8.7}$$

Term	Description of cost	Eq.
\mathcal{C}_f	Decreases difference in magnetic force	8.3
\mathcal{C}_d	Decreases distance between magnet and pen	8.4
\mathcal{C}_α	Encourages close distance over large force	8.5
\mathcal{C}_l	Decreases lag to path contour	8.6
\mathcal{C}_c	Decreases distance to path contour	8.6
\mathcal{C}_θ	Magnet stays ahead of pen	8.7
$\mathcal{C}_{\dot{\theta}}$	Ensures smooth progress	8.7

TABLE 8.2: Summary of costs terms used in optimization.

8.3.3 Dynamics model

To phrase electromagnetic haptic guidance in the MPCC framework, we contribute a dynamics model $f(\mathbf{x}, \mathbf{u})$ describing the system dynamics given its states \mathbf{x} and inputs \mathbf{u} .

$$\begin{aligned} \dot{\mathbf{x}} &= f(\mathbf{x}, \mathbf{u}) \text{ with} \\ \mathbf{x} &= [\mathbf{p}_m, \dot{\mathbf{p}}_m, \alpha, \theta] \in \mathbb{R}^6 \text{ and } \mathbf{u} = [\ddot{\mathbf{p}}_m, \dot{\alpha}, \dot{\theta}] \in \mathbb{R}^4. \end{aligned} \quad (8.8)$$

The system state \mathbf{x} consists of the position of the electromagnet $\mathbf{p}_m \in \mathbb{R}^2$ and its velocity $\dot{\mathbf{p}}_m$, the magnet intensity α and the current path progress θ . The inputs to the system \mathbf{u} consist of the in-plane electromagnet accelerations $\ddot{\mathbf{p}}_m$, and velocities $\dot{\alpha}$ and $\dot{\theta}$ for magnet intensity and the spline progress respectively. Note that we empirically found that magnet accelerations yield smoother motion than using velocities. The system model is given by the non-linear ordinary differential equations using first and second derivatives as inputs:

$$\ddot{\mathbf{p}}_m = v_m, \quad \dot{\alpha} = v_\alpha \quad \text{and} \quad \dot{\theta} = v_\theta, \quad (8.9)$$

where $v_{(\cdot)}$ are the external inputs. The continuous dynamics model $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u})$ is discretized using a standard forward Euler approach: $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t)$ [87].

In our hardware implementation, we derive the sets of admissible states χ and inputs ζ empirically to conform to the physical hardware constraints of the linear stage (e. g. max x,y-position) and EM specifications (e. g. max voltage). These are used in the constrained optimization problem solved in Eq. 8.11. The pen position is propagated via a standard linear Kalman filter [87]. While not an accurate user model, it works well in practice since the states are recalculated at every timestep.

8.3.4 Optimization

We combine the cost terms (Table 8.2) to control the force and position of the actuator to form the final stage cost:

$$\begin{aligned}
 J_k = & w_f \mathcal{E}_f(\mathbf{p}_{m,k}, \mathbf{p}_{p,k}, \alpha_k, \boldsymbol{\theta}_k) + \\
 & w_d \mathcal{E}_d(\mathbf{p}_{m,k}, \mathbf{p}_{p,k}) + w_\alpha \mathcal{E}_\alpha(\alpha_k) + \\
 & w_l \mathcal{E}_l(\mathbf{p}_{p,k}, \boldsymbol{\theta}_k) + w_c \mathcal{E}_c(\mathbf{p}_{p,k}, \boldsymbol{\theta}_k) + \\
 & w_\theta \mathcal{E}_\theta(\boldsymbol{\theta}_k) + w_{\dot{\theta}} \mathcal{E}_{\dot{\theta}}(\dot{\boldsymbol{\theta}}_k),
 \end{aligned} \tag{8.10}$$

where the scalar weights $w_l, w_c, w_\theta, w_{\dot{\theta}}, w_f, w_d, w_\alpha > 0$ control the influence of the different cost terms. The values used in our experiments and applications can be found in the Implementation section. The system states and inputs are computed by solving the N -step finite horizon constrained non-linear optimization problem at time instance t .

The final objective therefore is:

$$\begin{aligned}
 \underset{\mathbf{x}, \mathbf{u}, \boldsymbol{\theta}}{\text{minimize}} \quad & \sum_{k=0}^N w_k (J_k + \mathbf{u}_k^T \mathbf{R} \mathbf{u}_k) & (8.11) \\
 \text{Subject to:} \quad & \mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k) & \text{(System Model)} \\
 & \mathbf{x}_0 = \hat{\mathbf{x}}(t) & \text{(Initial State)} \\
 & \boldsymbol{\theta}_0 = \hat{\boldsymbol{\theta}}(t) & \text{(Initial Progress)} \\
 & \boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k + \dot{\boldsymbol{\theta}}_k dt & \text{(Progress along path)} \\
 & 0 \leq \boldsymbol{\theta}_k \leq L & \text{(Path Length)} \\
 & \mathbf{x}_k \in \mathcal{X} & \text{(State Constraints)} \\
 & \mathbf{u}_k \in \zeta & \text{(Input Constraints)}
 \end{aligned}$$

Here k indicates the horizon stage and the additional weight w_k reduces over the horizon, so that the current timestep has more importance than later timesteps. $\mathbf{R} \in \mathbb{S}_+^{n_u}$ is a positive definite penalty matrix avoiding excessive use of the control inputs. In our implementation we use a horizon length of $N = 10$. Experimentally we found that this is sufficient to yield robust solutions to problem instances and longer horizons did not improve results, yet linearly increases computation time.

8.4 IMPLEMENTATION

In this section, we detail our electromagnetic force model used in our optimal control scheme as well as the implementation of our hardware prototype.

8.4.1 Electromagnetic force model

Our approach requires a model of the interaction between a variable-strength electromagnet (EM) and the permanent magnet in the stylus that is sufficiently accurate and can be evaluated in real time. Accurately modeling the non-linear EM field of the electromagnet's core is typically done through finite element analysis (FEA), which cannot be performed in real time. Similarly, precomputing the volumetric map of the EM field \mathbf{B}_m via FEA for all levels of electrical current is not computationally feasible. We therefore contribute a novel fast approximate yet accurate electromagnetic model that provides a good balance between speed and accuracy to enable haptic guidance in applications such as writing or sketching.

In general, we aim at finding the actuation force on the pen \mathbf{F}_p , which is given by integrating over the volume of the permanent magnet in the pen:

$$\mathbf{F}_p = \iiint \nabla (\mathbf{M}_p \cdot \mathbf{B}_m(\cdot)) dx dy dz, \quad (8.12)$$

where \mathbf{M}_p is the magnetization of pen magnet and $\mathbf{B}_m(\cdot)$ is the EM field evaluated at the pen position, which is too costly to evaluate in real time. Our model approximates this underlying physical phenomena, can be efficiently evaluated at every iteration of our optimization procedure and provides a very good fit to empirical data. In this section, we briefly discuss the most important aspects of our model, for a full derivation and analysis we refer readers to the Appendix 4.1.2.

We make the following two assumptions in our derivation: 1) the electromagnet and the permanent magnet can be approximated as dipoles (i. e. oriented point magnets), and 2) for the smaller dipole (the permanent magnet in the pen) the out-of-plane vector component is much larger than the in-plane counterpart. This allows us to use only the vertical component in the calculation of the force.

The first assumptions allows us to use the standard model by Yung et al. [285] to compute the force exerted by the electromagnet \mathbf{m}_e onto the pen \mathbf{m}_p (see Figure 8.6) as:

$$\mathbf{F}_p = \frac{3\mu_0}{4\pi r_p^5} \left[(\langle \mathbf{m}_p, \mathbf{r}_p \rangle) \mathbf{m}_e + (\langle \mathbf{m}_e, \mathbf{r}_p \rangle) \mathbf{m}_p + \right. \\ \left. (\langle \mathbf{m}_p, \mathbf{m}_e \rangle) \mathbf{r}_p - \frac{5(\langle \mathbf{m}_p, \mathbf{r}_p \rangle)(\langle \mathbf{m}_e, \mathbf{r}_p \rangle)}{r_p^2} \mathbf{r}_p \right], \quad (8.13)$$

where μ_0 is a constant (vacuum permeability $4\pi \cdot 10^{-7}$ [H/m]) and \mathbf{r}_p is the 3D vector between the centers of the electromagnet and pen dipoles. In contrast to FEA, this expression is analytic and differentiable, thus suited for iterative optimization.




FIGURE 8.6: Illustration of the model to compute the force \mathbf{F}_p on dipole \mathbf{m}_p due to dipole \mathbf{m}_e .

Figure 8.6 shows all quantities needed to compute the total magnetic force exerted on the pen. The expression does, however, lead to an actuation force \mathbf{F}_p that depends on the tilt of the pen. In pre-tests, we found that users can not perceive a difference in strength when tilting the pen in-place. We therefore leverage our second assumption, which reduces the EM model from 6 DOF to 3 DOF, to avoid this computation.

Based on the second assumption, we can retrieve the two vertical force vectors of the electromagnet \mathbf{m}_m and the pen $\tilde{\mathbf{m}}_p$. The vector between the two centers can now be computed as \mathbf{r}_p . We then project this vector onto the plane, yielding the final vector d between the pen tip and the in-plane projection of the actuator dipole. The total force acting on the pen (Eq. 8.13) can now be decomposed as:

$$\mathbf{F}_p = F_a \mathbf{e}_d + F_z \mathbf{e}_z. \quad (8.14)$$

Here $\mathbf{F}_a = F_a \mathbf{e}_d$ represents the in-plane quantity we seek to control, as it is the magnitude F_a of the force vector \mathbf{F}_a in the direction of a unit vector \mathbf{e}_d along d . F_z is the vertical force components which pulls the pen downwards. During our experiments there was no significant change in perceived friction when comparing the drawings with and without electromagnet (i. e. with or without F_z). For this reason we do not actively account for F_z in our optimization and only maintain

the in-plane force contribution (\mathbf{e}_d direction). The actuation force as function of pen-magnet separation is obtained as:

$$\mathbf{F}_a = \alpha F_0 \left(\frac{d \left(4 - \frac{d^2}{h^2} \right)}{h \left(1 + \frac{d^2}{h^2} \right)^{\frac{7}{2}}} \right) \mathbf{e}_d, \quad (8.15)$$

where $\alpha \in [0, 1]$ is a dimensionless scalar to control the desired strength of the force that should be felt by users, h is the center-to-center distance between both magnets projected on to the z-axis (Figure 8.4) and F_0 is a constant force parameter given by the expression,

$$F_0 = \frac{3 \mu_0 m_p m_m}{4 \pi h^4}. \quad (8.16)$$

The actuation force F_a is zero if the two magnets are aligned with one another ($d = 0$), it has a maximum $F_a^{max} = 0.9 F_0$ at $d = 0.39h$, and we can assume there is no more attraction for distances $d > 2h$. Note that the second assumption (only use in-plane component) lead only to a small approximation error. Compared to an angle dependent formulation (see Appendix A.1.1.2), a tilt of up to $\theta = 30^\circ$ leads to a max error in our model (Equation A.14) equivalent to shifting the distance d by ± 3 mm. This uncertainty in d is comparable with the in-plane positioning error (dispersion) of our hardware prototype. An angle dependent formulation of our model (i. e. 6 DOF) can be found in Appendix A.1.1.2 for future use in cases where the pen angle is tracked. This model remains valid for other hardware implementations involving a single moving electromagnet or can be easily extended onto a grid of fix electromagnets.

8.4.2 Hardware prototype

We have developed one possible hardware instance that utilizes our optimization scheme for an in-plane haptic guidance system (see Figure 8.2). Our system consists of 3 main components: 1) a moving platform that controls the 2D location of the electromagnetic actuator, 2) an input device such as a stylus, and 3) an output device such as a digital tablet or digitizer used in combination with a non-digital drawing surface.

8.4.2.1 Motion platform

The motion platform consists of a controllable electromagnet on a bi-axial linear stage directly underneath the drawing plane. The linear stage has two orthogonal 12 mm lead screws, which are driven by two 24 V, 4.0 A NEMA23 high-torque

stepper motors. We control the motors with a DQ542MA stepper driver and an Arduino UNO. As electromagnet, we use an Intertec ITS-MS-5030-12VDC magnet (5 cm diameter, 3 cm height, 12 V), controlled via pulse-width modulation. It can deliver up to 488 mN of lateral force at 11 W. We used FEM analysis to select this magnet from a range of commercially available magnets [51]. It provides a balance between power consumption, size, and force, i. e. it delivers a strong perceivable force while having a small footprint relative to our hardware.

To measure the positional dispersion of the motion platform, we moved the electromagnet at full strength ($\alpha = 1$) to 300 random locations and then always back to the center of the drawing surface. During these trials, a user held the pen upright and followed the magnet passively. By measuring the difference in target and actual position, we found that our system yields 2.8 mm (± 0.8 mm) of point dispersion. We believe this is sufficient for most applications and our experiments. One of the factors that contribute to this dispersion is the vanishing of the actuation force F_a as $d \rightarrow 0$. This can lead to the pen motion stopping slightly before it reaches the target.

8.4.2.2 Software

Our software runs on a standard PC (Intel Core i7-4770 CPU 4 cores at 3.40 GHz) in all our experiments. The solver is implemented in FORCES Pro [63], which produces efficient C-code. The following weight values are used for our control scheme (Equation 8.10):

w_l	w_c	w_θ	$w_{\dot{\theta}}$	w_f	w_d	w_α	w_v	w_m
1.5	1.5	10.	0.1	10.	0.05	7.	1.	1.

Due to the steepness of the electromagnetic force \mathbf{F}_p and the potentially fast pen motion, runtime and latency are crucial performance metrics. The optimization algorithm contributes to both, whereas latency is dominated by the hardware and I/O. The mean solve time for a problem instance is 7.4 ms (± 3.0 ms). This can be expected to be mostly constant since we do not manipulate the system state space and the only measured input comes from the pen. The hardware and overall system latency adds to the solve time. We use a high-speed camera (1000 fps) to establish the motion (pen) to motion (magnet) latency. This yields an approximate latency of ~ 10 ms. Given the combined latency envelope of ~ 20 ms, we did not experience any abrupt pen snapping in our experiments.

8.4.2.3 Input and Output Devices

Our primary input and output devices for our user experiments consist of a 3D printed ballpoint pen with a permanent ring magnet mounted in the shaft (see Figure 8.2) and

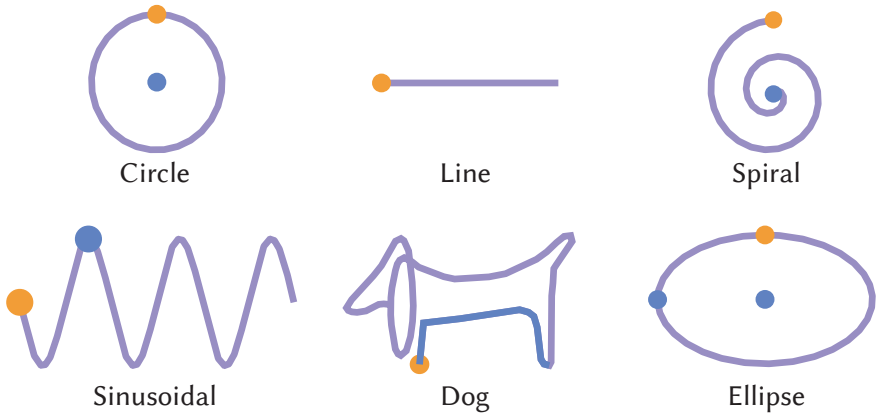


FIGURE 8.7: Shapes of our user tests. The drawing surface only contained sparse visual references (shown in blue) and starting positions (orange).

a piece of paper. The strokes are recorded by a Sensel Morph pressure sensitive touch pad [185]. If the system cannot locate the pen (e. g. when it is lifted) the last known position is used. We chose the Sensel for its high spatial resolution (6502 DPI), high speed (500 Hz) and low latency (2 ms), according to specification, and since the sketching surface does not interfere with the input recognition. Additionally, we developed an all digital input/output system with a multi-touch tablet + stylus. We use a Galaxy Tab A tablet with capacitive input and an off-the-shelf stylus with a permanent magnet placed near the tip. This magnet is slightly bigger (12 mm radius, 12 mm height) to compensate for the increase in tablet thickness.

8.5 EVALUATION

We first evaluated if our optimization scheme is beneficial for users in providing haptic guidance compared to a no-feedback baseline. In a second experiment, we compared our method with an open-loop and a closed-loop approach.

8.5.1 *Experiment 1 - Haptic feedback*

We compared our MPCC formulation with a no-feedback baseline to gather insights on task performance and user perception. Users were asked to draw several shapes (see Figure 8.7) and we evaluated accuracy and subjective feedback.

8.5.1.1 Procedure and tasks

We recruited 12 participants from the local university, all without professional drawing experience. Users were given an introduction to the system functionality and got to experience the system in a self-timed training phase. During the experiment we asked each participant to draw six basic shapes, each with and without haptic feedback. The presentation order of shapes and interface condition was counterbalanced. The drawing surface (white piece of paper) only contained a starting point and, in the case of more complex shapes, limited additional visual guidance (see Figure 8.7). Furthermore, the participants were shown a scaled version during task execution (scaled to prevent 1:1 copying). After the full experiment, users completed a questionnaire on their preference.

Quantitative Results

We compute the Hausdorff-like distance [224] between the drawn path and the reference path as error metric. To make the metric robust to drawing speed, we re-sample both paths equidistantly. To ensure fairness we also compute the inverse distance (reference to drawn path). A Kolmogorov-Smirnov test [133] indicated that the set of uni-directional distances is not significantly different from the set of inverse distances. We therefore only report uni-directional distances. The quantitative results for each target averaged over all participants are summarized in Table 8.3. Our method on average resulted in a 66% ($\pm 24.5\%$) lower error, i.e. it improve the average point-to-path difference by 1.54 mm. A two-way ANOVA on the mean error revealed a main effect for the feedback type ($F=46.187$, $p<.001$) and for the shapes ($F=11.771$, $p<.001$). Post-hoc analysis revealed that the line was statistically significantly different from all other shapes. This indicates that our method is beneficial for non-trivial shapes.

Qualitative Results

A brief exit interview shows that users subjectively rate the system favorably, on a 5-point Likert scale (5 is best), in terms of accuracy (4.33 ± 0.62), speed (4.00 ± 0.91), force (3.50 ± 0.86) and overall performance (4.50 ± 0.9). While we acknowledge that this might be in part due to novelty effects, we believe that the quantitative results indicate that our system is beneficial for users in general. The ratings indicate that participants generally see benefit in our approach and are not disturbed or hindered when using our approach.

Scenario	With		Without		Err %
	Mean	SD	Mean	SD	
Circle*	2.19	0.90	4.26	2.39	0.51
Line	1.18	0.80	1.03	0.84	1.15
Spiral*	2.55	0.75	4.38	1.64	0.58
Sinus*	2.53	0.70	5.08	2.19	0.50
Dog*	2.31	0.54	3.81	1.32	0.60
Ellipse*	2.40	0.56	3.84	1.22	0.62

TABLE 8.3: Mean accuracy (*mm*). * indicates statistical significance ($\alpha.05$).

8.5.2 Experiment 2: Comparison of control strategies

In this second experiment, we compared our time-free closed-loop optimization strategy to a simpler MPC variant and our implementation of dePEND [278], denoted as dePEND*.

8.5.2.1 Procedure and tasks

We asked twelve new participants (students and staff from a local university) to draw one complex shape (dog in Figure 8.7) in three different conditions: *dePEND**, time-dependent closed loop (*MPC*), and time-free closed loop (*Ours*), counterbalanced using a latin square. The speed of the system in the time-dependent cases was decided empirically based on pre-tests to work well at regular drawing speeds. After receiving instructions and a brief training, participants completed the three drawings. Participants were also encouraged to provide comments during the individual conditions.

8.5.2.2 Data collection

We analyze three measures: 1) the mean distance from the pen to the path, 2) the mean distance from the pen position projected onto the path and the setpoint along the path, denoted as $d(\text{pen}, \hat{s}(\cdot))$, and 3) the mean distance from the pen to the electromagnet. By taking the mean of the error terms over subjects we ensured equal numbers of datapoints, accounting for differences in speed. Participants were instructed to draw at roughly constant speed. This was done to ensure fairness in comparing our approach with the open-loop approach, which would deteriorate

	pen-path	d(pen, s(θ))	pen-empl
<i>dePENd*</i>	4.1(± 0.7)	38.0(± 56.9)	38.2(± 25.1)
MPC	3.9(± 1.3)	45.0(± 50.8)	8.6(± 1.6)
<i>Ours</i>	2.0 (± 0.6)	6.2 (± 0.8)	4.6 (± 0.9)

TABLE 8.4: Mean distances in *mm* for Experiment 2).

if the variability of the drawing speed were to high. Note that our approach does generally not require this assumption.

8.5.2.3 Quantitative results

Table 8.4 summarizes our quantitative findings. Not surprisingly, the distance from the electromagnet to the pen and $d(\text{pen}, s(\hat{\theta}))$ for *dePENd** is quite large. Since the force exerted on the pen falls-off quadratically with distance, participants often lost any haptic guidance early on, confirmed via user comments such as “I don’t feel anything” (P3) and “Is the system on?” (P6).

A Kruskal-Wallis test revealed that our approach has the highest accuracy compared to *dePENd** and MPC ($H(2)=20.76$, $p < .001$). Furthermore, the setpoint $s(\theta)$ ($H(2)=7.362$, $p < .05$) and the electromagnet ($H(2)=27.12$, $p < .001$) are closest to the pen using our approach. Thus our time-free formulation overcomes both problems of wrong setpoints (*MPC*) and a run-away electromagnet (*dePENd**). Figure 8.8 shows one typical example of a user. Both the distance along the path and the pen-magnet distance fluctuate strongly for *dePENd** and *MPC* control strategies. Our approach yielded a continuously low error.

While *MPC* reduces the distance from the pen to the magnet, it does not optimize for the progress along the path and hence may pull the pen into undesired directions. Furthermore, we saw that *MPC* produced extreme corner cutting behavior to catch up to the advancing setpoint. Both *dePENd** and *MPC* also produce results with high standard deviations. This is likely due to the absence of direct coupling between user feedback and path progress, which makes it possible for the user to lag behind the setpoint significantly (albeit at the cost of reduced force feedback). In our approach, the path progress is adjusted to the user’s drawing speed, drastically reducing the standard deviation and in consequence ensuring delivery of force feedback throughout the drawn path.




FIGURE 8.8: Comparison of error (path distance pen- $s(\theta)$) over time for a single participant (P1). The inverse u-shape illustrates that the setpoint $s(\theta)$ moves at a different speed than the user for *dePEND** and MPC. The data is smoothed to increase readability.

8.5.2.4 Qualitative results

From our observations we saw that $s(\theta)$ was either in front or behind the user for MPC. This was also confirmed in our interview, where people especially commented on the MPC strategy: “The system tries to push me in the wrong direction” (P2) and “It is counteracting me” (P11). In contrast with our formulation the magnet remains always slightly ahead of the pen, resulting in users commenting on our approach as the most preferred condition. In the words of one subject this is: “since I still had control” (P9).

In summary, these initial results indicate that our approach outperforms existing open-loop and time-dependent closed-loop approaches. *dePEND** causes numerous problems, including users not perceiving any haptic feedback. This is especially troublesome in settings where autonomy is desired. In MPC the haptic feedback is perceived, but can be erroneous. This is especially evident when users do not conform to the expected behavior. We plan to perform more in-depth experiments to investigate, for which applications our approach can be especially beneficial, and for which levels of autonomy.

8.6 APPLICATIONS

To further demonstrate the potential of our approach we illustrate possible usage-scenarios including calligraphy, outlining and inking. Finally, we combine the haptic




FIGURE 8.9: Our approach can be used as a guidance system for calligraphy, where users either follow a target path very closely, or deviate if desired.

feedback mechanism with a simple digital drawing application to initially explore the possibility of dynamic references.

Calligraphy

Figure 8.9 illustrates writing of flourished characters, with only minimal visual guidance (single starting point). Our system takes the character as input, users can then draw at their desired speed. Although an offset from the reference path remains, the lines are smooth and the overall shape is close to the desired characters.

Outlining & inking

Figure 8.10 illustrates the effect of two core capabilities of the proposed approach. Here we first outline the proportions of the dragon head (gray guidance lines) and then use different pens to ink-in the details. Note that the system provides haptic guidance but allows the user to draw the shape in different styles (e. g. the ears of the two upper dragons) and with varying high-frequency detail, while maintaining similarity to the reference shape. This is a direct consequence of using time-free closed loop control approach. In this case, all four variants were drawn without changes to the system or desired path.

Virtual tools

Using a digital tablet with capacitive display (Figure 8.11) we explore integrating dynamically changing references. In a sketching application, artist select different virtual tools, and position and configure these anywhere. The canvas and the haptic feedback system then pull the stylus towards these virtual guides. In Figure 8.11,




FIGURE 8.10: Different variants of the same dragon, drawn with identical system settings by a novice. Each pair of drawings used with different tools. First a pencil for proportions and a fine-liner (top) or pencil (bottom) to ink-in details. Multi-stroke lines are achieved by approaching each separate instance as a new drawing.



FIGURE 8.11: Virtual tools can be used to dynamically construct a reference path combining haptic and visual feedback. Here a simple drawing application combines freeform sketching with different virtual rules and guides that can be felt by the user.

the user has selected a tool that helps them when drawing an ellipse that snaps to a previous part of the drawing, both visually and in terms of haptics.

8.7 DISCUSSION

Our experiments indicate that the proposed approach indeed increases accuracy in drawing tasks and that users perceive the system favorably. While our system increased users' accuracy for complex shapes, it did not yield any improvements for the straight line. This limitation can be attributed to the maximum speed of the linear stage, as indicated by user feedback. In our interviews, some users indicated that they had the feeling that their drawings without feedback were more accurate once they experienced the haptic guidance. This suggests the potential for short-term muscle memory development when using our haptic guidance system. Long-term learning is a very interesting area to explore for our approach and haptic guidance systems in general. We plan to conduct such experiments in the future.

Our experiments currently focused on drawing and sketching applications. We believe our control strategy can be beneficial for a wide range of applications. We started exploring the usage of our approach with a tablet and digital stylus. Further experiments are necessary to determine the levels of complexity at which our approach is most beneficial. Allowing users to adjust their input on-demand is crucial, particularly since systems typically lack complete knowledge of the users' target paths. Our approach is a first step in the direction of balancing user input and system control for haptic guidance systems, and can be extended to other devices beyond electromagnetic systems if appropriate force models are provided.

In terms of hardware, the speed of the linear stage was a limiting factor. Future improvements may involve faster, more compact linear stages or a matrix of stationary electromagnets. The former requires no changes to our formulation, while the latter necessitates modifications to the EM model and dynamics model. A matrix of electromagnets could pave the way for a thinner form-factor design, which is an exciting research direction.

Addressing the hardware-induced speed limitation will open new avenues for efficient closed-loop control strategies, since faster pen motion would also tighten the latency and accuracy budget. Incorporating a mechanism to reconstruct the tilt of the pen would enhance sensing capabilities. This could be achieved for example via an accelerometer built into the pen or via a grid of hall-sensors underneath the surface. Information on the pen tilt could then be combined with the angle dependent formulation of our EM model. Furthermore, we believe there are many research opportunities in combining our approach with ink beautification approaches (e. g. [247, 248, 277]). Particularly interesting would be to leverage fully predictive models for non-drawing applications (e. g. DeepWriting [7]).

Future work could explore combining our approach with various types of haptic feedback, either environment mounted or body-worn, and different form factors such

as spherical electromagnets [145, 286]. Electromagnetic feedback in combination with spatial actuation maybe interesting in other settings. For example, a magnet mounted to a robotic arm could deliver contact-less feedback in VR scenarios. It would also be interesting to investigate how to best exploit the system capabilities in the context of motor memory and learning. All these scenarios make it necessary that a system interactively reacts to user input. Our approach enables such applications, and can generalize to such systems that go beyond 2D haptic guidance systems.

Finally, our optimization is subject to system dynamics, including user influence on pen position. Explicitly incorporating user behavior into system dynamics is challenging due to its non-linear nature. In the next chapter, we will explore a model-free learning-based approach that can learn underlying task structures and human behavior without explicit system dynamics.

We have proposed a novel optimization scheme for electromagnetic haptic guidance systems based on the MPCC framework. Our approach strikes a balance between user input and system control, allowing users to adjust their trajectory and speed on-demand. It optimizes system states and inputs over a receding horizon by solving a stochastic optimal control problem at each timestep. Our formulation provides dynamically adjustable forces and automatically regulates magnet position and strength. It can be evaluated analytically and is hence suitable for iterative, real-time optimization approaches. We implemented our approach on a prototype hardware platform and experimentally demonstrated that the feedback is well-received by users and offers higher accuracy compared to open-loop and time-dependent closed-loop approaches. We believe our approach offers a principled method for haptic guidance, enabling users to retain agency while receiving unobtrusive assistance. This approach has broad applications in areas such as drawing, sketching, writing, and guidance via virtual haptic tools.

MULTI-AGENT REINFORCEMENT LEARNING FOR GOAL-AGNOSTIC ADAPTIVE USER INTERFACES

In the previous chapter, we demonstrated the use of model predictive control to determine the actuation of a haptic feedback device. However, it relied on known system dynamics, including predictive models of user behavior. One way to overcome this limitation is through learned system, task, and user dynamics. In this chapter we investigate model-free reinforcement learning for such purpose. Furthermore, we switch away from haptics as use case and focus on adaptive point-and-click interfaces. This switch allows us to focus on the control, as the actuation and sensing in graphical user interfaces are deterministic. A core challenge in developing adaptive interfaces is inferring user intent and choosing adaptations accordingly. Current methods often depend on tediously hand-crafted rules or extensively gathered user data. Furthermore, heuristics need to be recrafted and data regathered for every new task and interface. To address this issue, we formulate interface adaptation as a multi-agent reinforcement learning problem. Our approach learns adaptation policies without relying on heuristics or real user data, enabling the development of adaptive interfaces across various tasks with minimal adjustments. In our formulation, a user agent mimics a real user and learns to interact with an interface via point-and-click actions. Simultaneously, an interface agent learns interface adaptations, to maximize the user agent's efficiency, by observing the user agent's behavior. For evaluation, we replaced the simulated user agent with actual users. Our study involved twelve participants and focused on automatic toolbar item assignment. Results demonstrate that the policies developed in simulation effectively apply to real users. These users were able to complete tasks with fewer actions and in similar times compared to methods trained with real data. "Additionally, we showcased the method's efficiency and generalizability across four different interfaces and tasks.

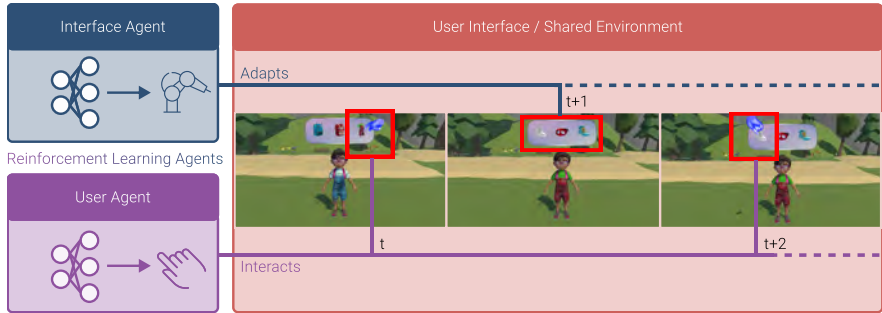


FIGURE 9.1: We formulate online user interface adaptation as a multi-agent reinforcement learning problem. Our approach comprises a user- and interface agent. The user agent interacts with an application in order to reach a goal and the interface agent learns to assist it. In the depicted example the user agent interacts with a Virtual Reality toolbar, while the interface agent assigns relevant items for the user agent. The interface agent does not know the goal of the user agent. Crucially, our approach does not rely on labeled offline data or application-specific handcrafted heuristics.

9.1 INTRODUCTION

In the previous chapter, we demonstrated the use of model predictive control to determine the actuation of a haptic feedback device, relying on known system dynamics that included predictive models of user behavior. However, these models have limitations. To overcome this, we investigate a model-free reinforcement learning approach for learning system, task, and user dynamics. In this chapter, we shift our focus from haptics to adaptive point-and-click interfaces, allowing us to emphasize control aspects, as actuation and sensing in graphical user interfaces are deterministic.

Point-and-click interfaces form the core interaction paradigm in modern Human-Computer Interaction (HCI) [124, 139, 220]. These interactions include clicking toolbar items, navigating hierarchical menus, and selecting objects—from UIs on traditional 2D desktops to emerging spatial UIs in immersive Mixed Reality. While individual point-and-click interactions are simple, complexity increases with the number of available interactive items. More items require users to process more information and evaluate more options, increasing cognitive load and diminishing the user experience, potentially extending task completion times.

Adaptive User Interfaces (AUIs) aim to mitigate this complexity by dynamically adjusting the interface to display the most relevant items for the user’s current task

while omitting less relevant information. This reduces cognitive load and simplifies navigation through complex interfaces, such as hierarchical menus.

Developing effective AUIs poses numerous challenges: AUIs must (1) infer users' intentions by observing their behavior and input, and (2) adapt the user interface based on these inferred intentions [204]. Previous research has predominantly relied on machine learning (ML) techniques [82, 238, 239, 266], or heuristics [33, 250, 254]. However, these methods depend on manually collected data or carefully hand-crafted rules, which is a significant limitation. Each new user interface or task requires new data collection and updated or redesigned adaptations, even when the interaction paradigm remains unchanged.

For example, customizing a game character through a toolbar interface illustrates these challenges Figure 9.1. The toolbar has limited slots, but the set of clothing items is much larger. An AUI must assign the most relevant items to the slots based on the user's previous selections and likely final outfits. Similarly, in a VR game where a user builds a castle from different blocks, the AUI recommends the most likely block to be selected for each operation, considering the user's past choices and probable final designs. Despite the similar point-and-click paradigm, ML or heuristic approaches would require separate data collection and rule design for each task. This requirement makes developing more complex AUIs time-consuming and costly.

In this paper, we propose *MARLUI*, a proof-of-concept framework for easy development of AUIs across various point-and-click UIs. We frame the point-and-click interaction paradigm in a Multi-Agent Reinforcement Learning (MARL) framework. An interface agent learns to adapt a UI by selecting the most relevant subset of items at each interaction step. The interface agent observes the user agent and minimizes its task completion time, while the user agent learns human-like point-and-click behavior to interact with the UI and train the interface agent. Unlike previous approaches requiring real user data to train an interface agent, our method relies on a human-like user agent. The user agent and the interface agent jointly learn by exploring the interface through trial-and-error. Our interface agent then transfers the learned adaptation strategy to real scenarios, selecting the most relevant items after a real user's click or selection. Switching tasks within the *MARLUI* framework, such as from dressing a game character to constructing a block castle, requires minimal modification and only training on the respective interface. This approach eliminates the need for developers to gather or create task- and interface-specific data or heuristics, streamlining AUI development. Future work can focus on user personalization, extending to more complex interfaces, adapting to changing goals, and addressing different tasks.

To model the point-and-click paradigm as a MARL problem, we propose a simulated user agent to learn interaction with a UI for task completion. We model the user agent hierarchically, decomposing decision-making, such as selecting a toolbar item, from motor control actions, such as moving the cursor to a desired menu slot. Targeting human-like behavior in the user agent, we incorporate cognitive and motor control bounds, following advancements in Computational Rationality [203]. Concurrently, we train an interface agent. The user agent aims to reach a goal state, while the interface agent adapts the UI to help the user agent achieve its goal more efficiently, minimizing the number of clicks. The user agent and interface agent operate in a turn-based manner: the user agent acts, followed by the interface agent’s adaptation, continuing until the task is completed. This novel approach enables our data- and heuristic-free method.

To demonstrate our framework’s effectiveness across various tasks and interfaces, we explore five use cases: (1) customizing a game character using a toolbar, (2) adaptive numeric keypad design, (3) building a block tower, (4) selecting initially out-of-reach objects, and (5) photo editing with anticipatory menu opening. These use cases cover a broad range of interface types, from heads-up displays to world-anchored interfaces, showcasing the versatility of our approach.

We evaluate our framework in two stages: in-silico studies on the character creation task, showing our interface agent can generalize to unseen goals, and evaluations with real participants, comparing *MARLUI* against data-driven baselines. Our findings indicate that training the interface agent with our simulated user agent enables real users to reduce the number of actions needed compared to previous frameworks.

In summary, we make three key contributions in this paper:

1. The first attempt at a multi-agent Reinforcement Learning approach that does not rely on real user data or hand-crafted heuristics to adapt point-and-click user interfaces in real-time.
2. A user agent that learns to operate a user interface and enables an interface agent to learn the task’s underlying structure purely by observing the user agent’s actions.
3. An interface agent that learns the task’s underlying structure by observing actions in the interface by the user agent.
4. Results from evaluations demonstrating the effectiveness of our approach and five use cases showcasing its general applicability for point-and-click tasks.

9.2 FRAMEWORK OVERVIEW

We define "adaptation" as the online alteration of an interface given current and past user input to support users in completing their task more efficiently. As such AUIs dynamically assign relevant interactive elements to an interface, thereby decreasing number of decisions and actions users need to take. MARLUI is a framework to create AUIs and can learn adaptive policies, that is neural networks that given an input predict suitable adaptations. MARLUI learns without needing data and only necessitates minimal adjustments to learn these policies for different AUIs. We specifically, assume the users to be Computational Rational [203] to make the problem tractable. In this section, we provide a high-level overview of the workings of our proposed framework. As shown in Fig. 9.2, MARLUI consists of three main components: a user interface as a shared environment (which gets adapted), a user agent (which learns to interact with the AUI), and an interface agent that learns an adaptation policy. Referencing the game character customization example illustrated in Fig. 9.1, we first describe the function of each component individually, followed by an explanation of how MARLUI can support real users.

USER INTERFACE / SHARED ENVIRONMENT Our approach models UI adaptation as a MARL problem, where an interface agent learns from a simulated user agent's interactions with a user interface, i.e., the shared environment. In this shared environment, user- and interface agent take actions in a turn-based manner. The user agent performs a point-and-click maneuver, which changes the state of the UI, e.g., altering the clothing of the game character. The interface agent observes the action and the change to the interface and based on that selects a new subset of items to display, e.g., adjusting the clothing items displayed in the toolbar. In turn, the user agent observes this adaptation and takes a new action. This cycle continues till the task is complete, e.g., the desired clothing configuration is reached.

USER AGENT The user agent aims to achieve its goal as fast as possible. In the context of game character creation, goals might involve selecting specific attributes for a character, such as the color of a shirt or backpack. The user agent can observe the visible parts of the interface, and knows its internal state. Based on these information, the user agent acts on the interface with point-and-click maneuvers, aiming to align the current state of the UI with its goal, e.g., aligning the current with the desired clothing configuration. By constraining its motor behavior in a physiologically plausible manner, the user agent is bound to exhibit human-like behavior for pointing and selecting items. Through trial and error, the user agent will eventually learn a policy that allows it to realize this alignment.

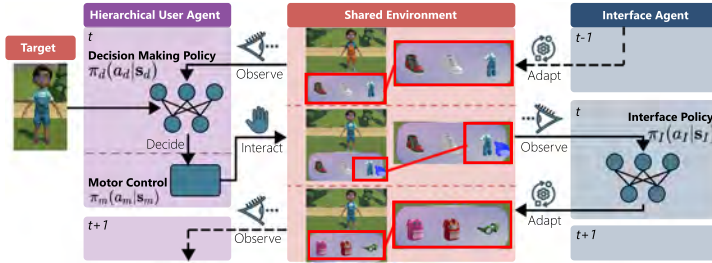


FIGURE 9.2: Our interface agent and user agent act in the same environment. The user agent is modeled as a two-level hierarchy with a high-level decision-making policy π_d and a low-level motor control policy π_m . The user agent interacts with the UI. The high-level policy of the user agent observes that state of the environment (Eq. 9.1) and chooses a specific menu slot as target accordingly (Eq. 9.2). The lower level receives this action and computes a movement (Sec. 9.3.2.2). The interface agent policy π_I adapts the interface to assist the user agent in achieving its task more efficiently. It observes user actions in the UI (Eq. 9.6) and decides on adaptations. Note that the interface agent cannot access the goal of the user, making the problem partially observable.

INTERFACE AGENT The interface agent adapts the UI in a turn-based manner to minimize the number of actions the user agent must perform to complete a task. Despite not knowing the user agent’s specific goal, it learns the task structure (i.e., sequence of states) through observing the user agent’s interactions with the UI. It can then learn to select the subset of items that are most relevant to the user agent at its current state, e.g., dynamically populating the toolbar with the clothing items that are most likely to be picked. Through trial and error, it can learn UI adaptation policies without relying on pre-collected user data or predefined heuristics.

INTERACTION WITH REAL USERS By mimicking human-like point-and-click behavior through the user agent, the interface agent can learn to adapt UIs such that it also assist real users in accomplishing the same task. To apply the learned adaptive interface to real users, the setting is changed such that interface agent interacts with the actual users instead of the user agent. In a turn-based fashion, it selects the most relevant next items after each click or selection of the actual user according to what it has learned through interactions with the user agent.

9.3 METHOD

We first present an outline the model of our user agent, consisting of a high- and low-level policy. Then we present the interface agent (Fig. 9.2).

9.3.1 *General Task Description*

We model tasks to be completed if the user agent achieves their desired goal. For game character creation, a goal can be the desired configuration of a character with a certain shirt (red, green, blue), and backpack (pink, red, blue). We represent the goal as a one-hot vector encoding \mathbf{g} of these attributes. A one-hot vector can be denoted as \mathbb{Z}_2^j , where j is the number of items. For the previous example, \mathbf{g} would be in \mathbb{Z}_2^6 as it possesses six distinct items.

Furthermore, the user agent can access an input observation denoted by \mathbf{x} , e.g., this can correspond to the current character configuration. The current input observation, \mathbf{x} , and the goal state \mathbf{g} are identical in dimension and type.

The user agent interacts with the interface and attempts to match the input observation and goal state as fast as possible, such that $\mathbf{x} = \mathbf{g}$. Each interaction updates \mathbf{x} accordingly, and a trial terminates once they are identical. In the character creation example, this would be the case if the shirt and backpack of the edited character are the same as the desired configuration. The interface agent makes online adaptations to the interface. It does *not* know the specific goal of a user. Instead, it needs to observe user interactions with the interface to learn the underlying task structure that will yield the optimal adaptations, e.g., the user likely wants to configure the shoes after configuring the backpack.

The interface agent learns to adapt the UI to the user agent by maximizing the same expected discounted reward. Specifically, the interface agent learns to infer optimal next adaptations over an infinite horizon by updating implicit probabilities of likely next actions of the user agent, given its current sequence of past actions. As such it does not learn an explicit or implicit goal probability distribution, but a distribution of the most likely next actions of the user agent. An example is to suggest a white and blue shoe to the user agent in the tool, as the user agent has not interacted with the category of shoes thus far (Fig. ??).

9.3.2 *User Agent*

The user agent interacts with an environment to achieve a certain goal (e.g., select the intended attributes of a character). The agent tries to accomplish this as fast and accurately as possible, hence it minimizes task completion time. Thus, the

user agent first has to compare the goal state and input observation and then plan movements to reach the target. We model the user agent as a hierarchical agent with separate policies for a two-level hierarchy [**Langerak:2021:Generalizing**]. First, we introduce a high-level decision-making policy π_d that computes a target for the agent (high-level decision-making), we approximate visual cost with the help of existing literature [50]. Second, a WHO Model Fitts’-Law-based low-level motor policy π_m that decides on a strategy to reach this target. We now explain both policies in more detail.

9.3.2.1 High-level Decision-Making Policy

The high-level decision-making policy of the hierarchy is responsible to select the next target item in the interface. The overall goal of the policy is to complete a given task while being as fast as possible. Its actions are based on the current observation of the interface, the goal state, and the agent’s current state. More specifically, the high-level state space S_d is defined as:

$$S_d = (\mathbf{p}, \mathbf{m}, \mathbf{x}, \mathbf{g}), \quad (9.1)$$

which comprises: i) the current position of the user agent’s end-effector normalized by the size of the UI, $\mathbf{p} \in I^n$ (where n denotes the dimensions, e.g., 2D vs 3D), ii) an encoding of the assignment of each item $\mathbf{m} \in \mathbb{Z}_2^{n_i \times n_s}$, with n_i and n_s being the number of menu items and environment locations, respectively, iii) the current input state $\mathbf{x} \in \mathbb{Z}_2^{n_i}$, and iv) the goal state $\mathbf{g} \in \mathbb{Z}_2^{n_i}$. Here, I denotes the unit interval $[0, 1]$, and \mathbb{Z}_2^b is the previously described set of integers. The item-location encoding m represents the current state of a UI. It can be used, for instance, to model item-to-slot assignments. The action space A_D is defined as:

$$A_d = \mathbf{t}, \quad (9.2)$$

which indicates the next target slot $\mathbf{t} \in \mathbb{N}_{n_s}$. The reward for the high-level decision-making policy consists of two weighted terms to trade-off between task completion accuracy and task completion time: i) how different the current input observation \mathbf{x} is from the goal state \mathbf{g} , and ii) the time it takes to execute an action. Therefore, the high-level policy needs to learn how items correlate with the task goal as well as how to interact with any given interface. With this, we define the reward as follows:

$$R_d = \underbrace{\alpha \mathcal{E}_{gd}}_i - (1 - \alpha) \underbrace{(T_D + T_M)}_{ii} + \mathbb{1}_{\text{success}}, \quad (9.3)$$

where \mathcal{E}_{gd} is the difference between the input observation and the goal state, α a weight term, T_M the movement time as an output of the low-level policy, T_D

the decision time, and $\mathbb{1}_{\text{success}}$ an indicator function that is 1 if the task has been successfully completed and 0 otherwise.

In addition to movement time, we also need to determine the decision time T_D . To this end, we are inspired by the SDP model [50]. This model interpolates between an approximated linear visual search-time component and the Hick-Hyman decision time [98], both functions take into account the number of menu items and user parameters.

We define the difference \mathcal{E}_{gd} between the input observation \mathbf{x} and the goal state \mathbf{g} as the number of mismatched attributes:

$$\mathcal{E}_{gd} = - \sum_{x \in \mathbf{g}, y \in \mathbf{x}} \frac{\mathbb{1}_{x \neq y}}{n_{attr}}, \quad (9.4)$$

where $\mathbb{1}$ is an indicator function that is 1 if $x \neq y$ and else 0, x and y are individual entries in the vectors \mathbf{g} and \mathbf{x} respectively, and n_{attr} is the number of attributes (e.g., shirt, backpack, and glasses).

9.3.2.2 Low-Level Motor Control Policy

The low-level motor control policy is a non-learned controller for the end-effector movement. In particular, given a target, it selects the parameters of an endpoint distribution (mean $\mu_{\mathbf{p}}$ and standard deviation $\sigma_{\mathbf{p}}$). We set $\mu_{\mathbf{p}}$ to the center of the target. The target \mathbf{t} is the action of the higher-level decision-making policy (A_D). Following empirical results [74], we set $\sigma_{\mathbf{p}}$ to 1/6th of a menu slot width to reach a hitrate of 96%.

Given the current position and the endpoint parameters (mean and standard deviation), we compute the predicted movement time using the Fitts' Law derived WHO Model [93].

$$T_M = \left(\frac{k}{(\sigma_{\mathbf{p}}/d_{\mathbf{p}} - y_0)^{1-\beta}} \right)^{1/\beta} + T_M^{(0)}, \quad (9.5)$$

where k and β are parameters that describe a group of users, $T_M^{(0)}$ is the minimal movement time, and y_0 is equal to the minimum standard deviation. The term $d_{\mathbf{p}}$ indicates the traveled distance from the current position to the new target position $\mu_{\mathbf{p}}$. We follow literature for the values of other parameters [93, 118]. We sample a new position from a normal distribution: $\mathbf{p} \sim \mathcal{N}(\mu_{\mathbf{p}}, \sigma_{\mathbf{p}})$.

9.3.3 Interface Agent

The interface agent makes discrete changes to the UI to maximize the performance of the user agent. In our running example of character customization, it assigns

items to a toolbar to simplify their selection for the user agent. Unlike the user agent, we model the interface agent as a flat RL policy. The state space S_I of the interface agent is defined as:

$$S_I = (\mathbf{p}, \mathbf{x}, \mathbf{m}, \mathbf{o}), \quad (9.6)$$

which includes: i) the position of the user $\mathbf{p} \in I^2$, ii) the input observation $\mathbf{x} \in \mathbb{Z}_2^{n_i}$, iii) the current state of the UI $\mathbf{m} \in \mathbb{Z}_2^{n_i \times n_s}$, and iv) a vector including the history of interface elements the user agent interacted with (commonly referred to as stacking). The action space $A_I \in \mathbb{Z}$ and its dimensionality is application-specific. The goal of the interface agent is to support the user agent. Therefore, the reward of the interface agent is directly coupled to the performance of the user agent. We define the reward of the interface agent to be the reward of the user agent’s high-level policy:

$$R_I = R_D. \quad (9.7)$$

Note that the interface agent does *not* have access to the user agent’s goal \mathbf{g} or target \mathbf{t} . To accomplish its task, the interface agent needs to learn to help the user agent based on an implicit understanding of i) the objective of the user agent, and ii) the underlying task structure. Our setting allows the interface agent to gain this understanding solely by observing the changes in the interface as the result of the user agent’s actions. This makes the problem more challenging but also more realistic.

The interface agent learns an implicit distribution of possible goals, and by observing the user agent it narrows down the distribution over goals. At every time step the interface agent suggests the items that are most likely needed, given the goal distribution.

9.4 IMPLEMENTATION

We train the user and interface agent’s policies simultaneously in a shared environment (the AUI). All policies receive an independent reward, and the actions of the policies influence a shared environment. We execute actions in the following order: (1) the interface agent’s action, (2) the user agent’s high-level action, followed by (3) the user agent’s low-level motor action. The reward for the two learned policies is computed after the low-level motor action has been executed. The episode is terminated when the user agent has either completed the task or exceeded a time limit.

We implement our framework in Python 3.8 using RLLIB [154] and Gym [32]. We use PPO [231, 283] to train our policies. We use 3 cores on an Intel(R) Xeon(R) CPU @ 2.60GHz during training and an NVIDIA TITAN Xp GPU. Training takes

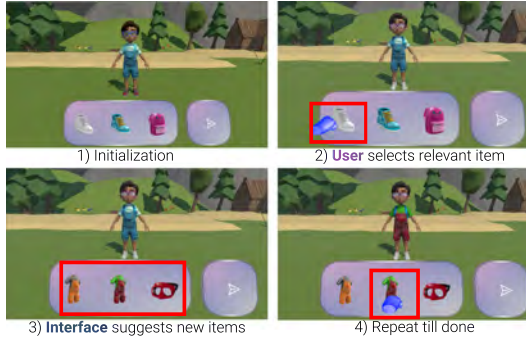


FIGURE 9.3: In our proposed task the interface agent and the user agent interact in a turn-based manner, in which the user (agent) matches a game character selection to a target state (1). First, the user operates a toolbar with three slots (2). Second, The interface agent assigns the most relevant items to the available slots (3). Finally, This cycle continues till the two characters match (4).

~ 36 hours. The user agent’s high-level decision-making policy π_u is a 3-layer MLP with 512 neurons per layer and ReLU activation functions. The interface agent’s policy π_i is a two-layer network with 256 neurons per layer and ReLU activation functions. We sample the full state initialization (including goal) from a uniform distribution. We use stochastic sampling for our exploration-exploitation trade-off.

We use curriculum learning to increase the task difficulty and improve learnability. Specifically, we adjust the difficulty level every time a criteria has been met by increasing the mean number of initial attribute differences. More initial attribute differences result in longer action sequences and are therefore more complex to learn. We increase the mean by 0.01 every time the successful completion rate is above 90% and the last level up was at least 10 epochs away.

We randomly sample the number of attribute differences from a normal distribution with standard deviation 1, normalize the sampled number into the range $[1, n_a]$ and round it to the nearest integer, where n_a is the number of attributes of a setting (in the case of game character $n_a = 5$).

9.5 EVALUATION

MARLUI aims to learn UI adaptations from simulated users that can support real users in the same task. Specifically, we want our framework to produce AUIs that are competitive with baselines that require carefully collected real user data. In this section, we evaluate if our approach achieves this goal. Thus, we first conduct an

in silico study to analyze how the interface agent and the user agent solve the UI adaptation problem in simulation. Then, we conduct a *user study* to investigate if policies of the interface agent that were learned in simulation benefit real users in terms of number of actions needed to complete a task.

9.5.1 Task & Environment

To conduct the evaluation, we introduce the *character-creation task* (see Fig. 9.3). In this task, a user creates a virtual reality game character by changing its attributes. A character has five distinct attributes with three items per attribute: i) shoes (red, blue, white), ii) shirt (orange, red, blue), iii) glasses (reading, goggles, diving), iv) backpack (pink, blue, red), and v) dance (hip hop, break, silly). The characters' attribute states are limited to one per attribute, i.e., the character cannot be dancing hip hop and break simultaneously. This leads to a total of 15 attribute items and 243 character configurations. We sample the goal uniformly from the different configurations.

The game character's attributes can be changed by selecting the corresponding items in a toolbar-like menu with three slots. The user can cycle through the items by selecting "Next." The static version of the interface has all items of an attribute assigned to the three slots, and every attribute has its own page (e.g., all shoes, if the user presses next, all backpacks). The character's attribute states correspond to the current input state \mathbf{x} and the target state \mathbf{g} , where \mathbf{g} is only known to the user agent. The goal of the interface agent is to reduce the number of clicks necessary to change an attribute, by assigning the relevant items to the available menu slots. For the *user agent*, the higher level selects a target slot, and the lower level moves to the corresponding location.

9.5.2 In Silico

TRAINING We evaluate the training of our framework against a static and a random interface. In the random interface, items are randomly assigned to the slots. Our method would perform on par, or worse, with random if it was incapable of learning. Hence, the random baseline provides a lower bound for performance. The static baseline has no adaptations, allowing us to evaluate the general effectiveness of our adaptive interface. For most tasks, a task-specific heuristic could be found that presents an upper baseline. However, this has to be designed specifically for each task. On the other hand, our goal is to provide a general framework that works across tasks with minimal changes to the reward function and observation spaces.

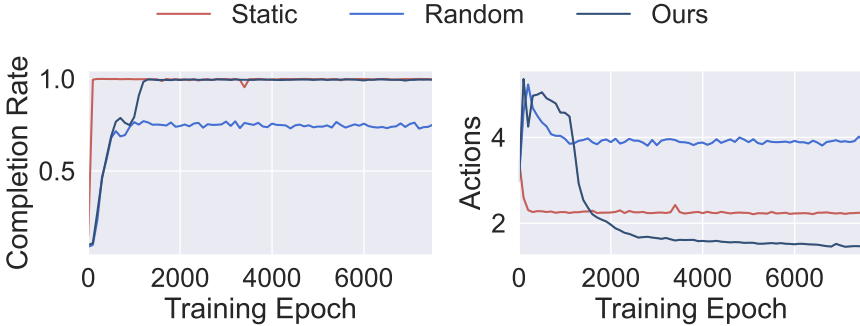


FIGURE 9.4: We train our agent till convergence. Left: the fraction of successfully completed episodes per epoch. Ours and Static reach a 100% successful completion rate. Random does not converge. Right: The number of actions needed on average during a successful episode. Our framework needs less actions compared to Static and Random.

Figure 9.4 shows the user agent’s task completion rate and number of actions per task of all three interfaces during training. Ours and the static baseline converge, whereas the random baseline does not. Furthermore, the mean number of actions of ours is lower than the mean of the static interface.

GENERALIZATION TO UNSEEN GOALS To understand how well our approach can generalize to unseen goals, we ablate the fraction of goals the agents have access to during training. We then evaluate the learned policies against the full set of goals, which is defined as all possible combinations of character attributes. The results are presented in Figure 9.5. We find that having access to roughly half of the goals is sufficient to not impact the results. This indicates that our approach generalizes to unseen goals of the same set for this task.

UNDERSTANDING POLICY BEHAVIOR We qualitatively analyzed the learned policies of our interface agent to understand how it supports the user agent in its task. In Figure 9.6, we show a snapshot of two sequences with identical initialization. To reach the target character configuration, the user agent can either select the blue bag or the purple glasses (both are needed). Depending on which item the user selects at this time step, the interface agent proposes different suggestions in subsequent steps. For instance, it might happen that the interface agent suggest two relevant items; while the user can only select a single one. As example, the blue backpack the user did not select initially (Figure 9.6, bottom) gets suggested again later; as the user




FIGURE 9.5: The fraction of successfully completed episodes as function of the fraction of the total number of goals. The graph shows that it is sufficient to see half of the goals to learn policies that generalize to all goals.

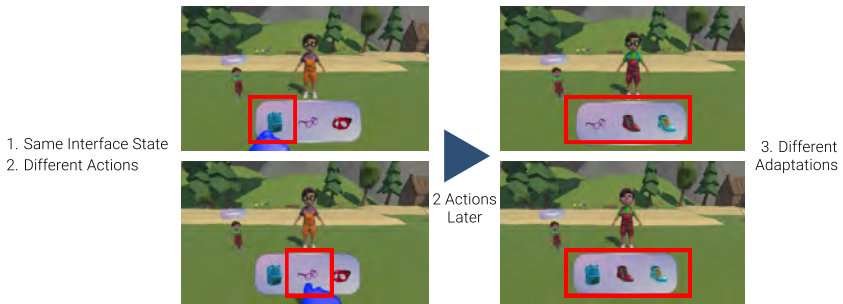


FIGURE 9.6: With our framework, multiple relevant items can be assigned simultaneously; yet the user can only select one (left). Depending on the user’s action (top: select backpack, bottom: select glasses), other item gets assigned later (top: shoes, bottom: backpack). Note that the user has not selected any backpack in those two actions. This shows that our framework actively adapts to user input.

has not selected a single backpack during the interaction cycle. This behavior shows that the interface agent implicitly reasons about the attribute that the user intends to select based on previous interactions. In short, the interface agent learns to suggest items that the user is not wearing, or that the user has not interacted with.

9.5.3 User Study

Our goal was to create a user agent whose behavior resembles that of real users, so the interface agent can support them in the same task. To this end, we evaluated the

sim-to-real transfer capabilities of our framework by conducting a user study where the interface agent interacted with participants instead of the user agent.

9.5.3.1 *Baselines*

We compared our framework to two supervised learning approaches and the static interface (see Sec. 9.5.1). In line with previous work [82], we used a Support Vector Machine (SVM) with a Radial basis function (RBF) kernel as a baseline. It represents a direct competitor to our approach, as it generalizes on a method level. However, it needs data recollection for every interface and task.

We used the implementations of scikit-learn [207] and optimized the hyperparameters for performance. The feature vector of the baseline was identical to that of our framework. The baseline learned the probability with which a user will select a certain character attribute next. We assigned the three attributes with the highest probability to the menu slots. Note that we did not consider "Next" to be an item.

DATASET We collected data from 6 participants to train the supervised baselines. These participants did not take part in the user study. They interacted with the static interface, which resulted in a dataset with over 3000 logged interactions. We saw that the SVM performance saturates when increasing to more than 2700 data points and reached around 91% top-3 classification accuracy (i.e., the percentage of how often the users' selected item was in the top three of the SVM output) on a test set. Furthermore, we found that the baseline generalize well to unseen participants (again through cross-validation). We used all data points for the final model.

METRICS We used two metrics (dependent variables) to evaluate our approach.

1. *Number of Actions*: the number of clicks a user needed to complete a task, which is a direct measure of user efficiency [37].
2. *Task Completion Time*: the total time a user needed to complete a task.

9.5.3.2 *Procedure*

Participants interacted with the interface agent and the two baselines. The three settings were counterbalanced with a Latin square design, and the participants completed 30 trials per setting. In each condition, we discarded the first six trials for training. The participants were instructed to solve the task as fast as possible while reducing the number of redundant actions. They were allowed to rest in-between trials. We ensured that the number of initial attribute differences (IAD), which refers to the number of clothing attributes that differ between the current selection and the




FIGURE 9.7: The average number of actions (left) and the task completion time (right) participants needed to finish the tasks of our user study, plotted with the 95% confidence interval. We compare Ours against Static, and an SVM. Our approach outperforms both baselines on the number of actions. Averaged over all participants and all trials.

target at the start of a trial, was uniformly distributed within the participant’s trials. Participants used an Oculus Quest 2 with its controller.

We recruited 12 participants from staff and students of an institution of higher education (10 male, 2 female, aged between 23 and 33). All participants were right-handed and had a normal or correct-to-normal vision. On average, they needed between 35 to 40 minutes to complete the study.

9.5.3.3 Results

We present a summary of our results in Figure 9.7. We analyzed the effect of conditions on the performance of participants with respect to the number of actions and task completion time.

Participants needed on average 3.34 actions to complete a task with our framework, compared to 5.73, and 3.87 for the static, and SVM baselines respectively. The normality and sphericity were not violated. We performed a repeated-measures ANOVA. We found a significant effect on method ($F(2,22) = 209.68, p < .001$). With a Holm-corrected post-hoc we find that both the SVM and Ours significantly outperform the static interface (both $p < .001$) in terms of number of actions. We also find that ours significantly outperforms the SVM ($p = 0.006$).

It is important to decompose the results per initial attribute difference. This is crucial because a task with a single IAD is expected to be shorter in duration than tasks with multiple IADs. We report the values in Table 9.1. Using a repeated-measures ANOVA, we found a significant effect ($F(8,0.25) = 35.537, p < .001$) on the interaction between the approach and the IAD. We limited our analysis to scenarios where the interaction involves "Ours" and only considered scenarios

	1	2	3	4	5
Heuristic	1.708 ± 0.396	3.889 ± 0.451	6.139 ± 0.797	7.903 ± 0.845	9.000 ± 0.000
SVM	1.456 ± 0.219	3.747 ± 0.515	4.607 ± 0.174	5.439 ± 0.259	5.827 ± 0.389
Our	1.625 ± 0.569	2.036 ± 0.297	3.281 ± 0.657	4.719 ± 0.613	5.311 ± 0.186

TABLE 9.1: User study result reported by IAD and method (Mean ± SD).

where the IADs between the methods were identical (i.e., we do not discuss, for example, the SVM with five initial attribute differences versus the static method with a single difference). Our findings from a Holm-corrected post-hoc analysis showed no difference between the static baseline and the other approach for a single initial attribute difference (all $p = 1.00$) and a significant difference for more than one IAD (all $p < .001$). When comparing the SVM versus our method, we found no difference for one IAD ($p=1.00$) or for five ($p=0.26$); for all other cases, a difference was found (all $p < .001$). We observed that the interface agent’s strategy, which groups similar category items (e.g., showcasing pink and blue backpacks if the character has a red one cf. Fig. 9.6), increases the chance of selecting a useful item. However, this strategy doesn’t apply when all attributes need changing.

When looking at the task completion time, participants using our framework needed 12.14 seconds to complete the task. The completion time was 12.36 for the static interface and 12.71 for the SVM. The task completion time was normally distributed (Shapiro-Wilk $p > 0.05$). We found no significant difference in overall task completion time with a Greenhouse-Geisser (for sphericity) corrected repeated-measures ANOVA ($F(1.568, 17.243) = 0.86, p = 0.42$). This could be, because despite requiring more action for a static baseline, the participants formed a strategy.

9.5.4 Discussion

To analyze if our multi-agent framework is competitive with a baseline that requires carefully collected real user data, we compared it against a supervised SVM. We did not find significant differences between the two approaches in the task performance metrics of the number of actions and task completion time. This suggests that our approach is a competitive alternative to data-driven approaches for creating adaptive user interfaces.

The adaptive approaches significantly reduce the number of actions necessary to complete the task compared to the static interface. However, no significant differences in task completion times were found. We argue that this could be due to real users being more familiar with the ordering of items in the static interface

that is kept constant across trials. This familiarity is not captured by our current cognitive model or incentivized in the reward function. In future work, we will model familiarity and investigate its effect on task completion time.

We have shown qualitatively that our interface agent learns to take previous user actions into account. This characteristic is core to meaningful adaptations. At the moment, our agent’s capabilities are limited by the size of the stack \mathbf{o} . In the future, recurrent approaches such as LSTM could be investigated to overcome this limitation.

Furthermore, we presented evidence that our framework can generalize to goals that were not seen during training. It is important to mention that the results of this experiment are subject to its task and that seen and unseen goals are from the same distribution. Nevertheless, the study provides first indications that our approach generalizes to applications where not all users’ goals might not always be encountered during training.

9.6 APPLICATIONS

To demonstrate the versatility of our framework, we introduce four additional point-and-click interfaces to demonstrate how our approach generalizes to different scenarios. Every scenario offers a distinct adaption, task, and interface. Our method requires minimal to no adaptations for the different scenarios (i.e., mostly a change in the dimensions of the observation and action spaces). We showcase both 2D interfaces as well as Mixed Reality interfaces. We show how the interface agent selects from a set of pre-designed UI widgets, hand user’s the correct blocks when building a tower, move out-of-reach items closer, and make hierarchical menus more efficient during photo editing.

Please refer to our supplementary video for visual demonstrations of the tasks. Due to the diverse nature of use cases, we will report either number of clicks or task completion time as success metric. All scenarios are showcased through the interaction of a real user with a trained interface agent. The numerical results are obtain in simulation.

9.6.1 *Number Entry*

We introduce a price entry task on a keypad. The interface agent selects a widget from a pre-designed set. We show that our approach can support applications requiring users to issue command sequences and provide meaningful help given a user’s progress in the task (see Fig. 9.8). The task assumes a setting where the simulated user must enter a product price between 10.00 and 99.99. To complete the task, the




FIGURE 9.8: Adaptive keypad: the user agent is asked to enter a randomly initialized price by using a keypad (1). The interface agent selects from three pre-designed different widgets either a normal keypad, a digits-only keypad or a non-digits-only keypad. The user agent selects a button of the chosen widget. The task ends when the user agent presses enter (4).

user agent has to enter the first two digits, the decimal point, the second two digits, and then press enter. The interface agent can select one of three different interface layouts: i) a standard keypad, ii) a keypad with only digits and iii) a widget with only the decimal point and the enter key.

The goal difference penalty (Eq. 9.3) in this case is based on whether the current price \mathbf{x} matches the target price \mathbf{g} :

$$\mathcal{E}_{gd} = - \sum_t \mathbb{1}_{\mathbf{x}_t \neq \mathbf{g}_t}, \quad (9.8)$$

where $\mathbb{1}$ is an indicator that is 1 if $\mathbf{x}_t \neq \mathbf{g}_t$ and 0 otherwise, and t is the current timestep. Every time a button is hit, t increases by 1. This is similar to the penalty in all other tasks. However, it considers that the order of the entries matters. On average, the user agent needs 4.0 seconds to complete the task in cooperation with the interface agent, compared to 4.9 seconds when using a static keypad. The number of clicks is identical, since the full task can be solved on the standard keypad.

Qualitative Policy Inspection

We observe that the interface agent learns to select the UI that has the biggest buttons for an expected number entry (e.g., only digits or only non-digits). From this we can conclude that the interface agent implicitly learns the concept of Fitt’s law and prioritizes larger buttons where appropriate.

9.6.2 *Block Building*

The second scenario is a block-building task (Fig. 9.9) where the user constructs various castle-like structures from blocks. Compared to the game character task,




FIGURE 9.9: Block Building: The user is building a castle from blocks (1). The user places the first block (2). The interface agent suggests a next block to place (3). This is repeated till the castle is built (4).

only the dimensionality of the observation and action space needs to be changed. It can choose between 4 blocks (wall, gate, tower, roof) and a delete button. The agent needs to move the hand to a staging place for the blocks (see Figure 9.9) and then place the block in the corresponding location. The block cannot be placed in the air, i.e., it always needs another block on the floor below. The interface agent suggests a next block every time the user places a block. However, the user can put the block down, in case it is unsuitable. An action is picking or placing a block.

This task represents a subset of tasks that do not have a Heads-Up-Display-like UI to interact with, but are situated directly in the virtual world. This is a common interactive experience of AR/VR systems. The user needs on average 1.1 actions with our framework, compared to 2.0 actions without the interface agent. Thus 1.1 indicates that the interface agent suggests the correct next block, most of the time.

Qualitative Policy Inspection

We observe that the policy learns to always suggest a block that is usable given the current state of the tower. This indicates that the policy has an implicit understanding of the order of blocks and can distinguish between those belonging to the foundation versus the upper parts of a tower.

9.6.3 *Out-of-reach Item Grabbing*

In the third usage scenario, the user needs to use their hand to grab an object that is initially out of reach. Thus, the interface agent needs to move an object within reach of the user, which can then grab it. The interface agent observes the location of the user's hand. The task environment includes several objects. Uniquely, in this




FIGURE 9.10: Out-of-reach object grabbing: the user agent attempts to grab a specific object, that is initially out of reach, in a space containing multiple objects (1). The user agent learned to move towards an object to indicate its intention to grab it (2). Based on that, the interface agent learned to move the intended object within the user agent’s reach (3). The user agent then grabs the object to finish the task (4).

scenario the user and the interface agent are forced to collaborate to select the correct target object and complete the task (see Figure 9.10); as it is impossible for the user agent to complete the task on its own. Compared to the game character task, only the dimensionality of the observation and action space needs to be changed.

In this use case, we changed the lower level of our user to learn motor control with RL instead of using the Fitts-Law-based motor controller. This highlights the modularity of our approach and can be useful in scenarios where existing models, such as Fitts’ Law, are not sufficient. The low-level motor control policy controls the hand movement. In particular, given a target slot, the policy selects the parameters of an endpoint distribution. Given the current position and the endpoint parameters (mean and standard deviation), we compute the predicted movement time using the WHO Model [93]. The low-level policy needs to learn i) the coordinates and dimensions of menu slots, ii) an optimal speed-accuracy trade-off given a target slot, and its current position. Refer to Appendix B.1 for more details.

Qualitative Policy Inspection

We find that the the policy selects objects positioned in the direction of the user’s arm movement rather than the closest ones. This indicates that the policy implicitly learns the correlation between directionality of movement and intent.



FIGURE 9.11: We introduce a photo editing task where (1) a user matches a photo to a target by operating a hierarchical menu. (2) The user selects the submenu ‘size’. (3) The user then selects the attribute ‘small’, which alters the image. (4) After the user has changed an attribute, the interface observes the new state of the photo and finds the most likely submenu for the next user action. (5) The user clicks on an item in the submenu to complete the task.

9.6.4 2D Hierarchical Menu

In this task, a user edits a photo by changing its attributes. A photo has five distinct attributes with three states per attribute: i) filter (color, sepia, gray), ii) text (none, Lorem, Ipsum), iii) sticker (none, unicorn, cactus), iv) size (small, medium large), and v) orientation (original, flipped horizontal, and vertical). The photo’s attribute states are limited to one per attribute, i.e., the photo cannot be in grayscale and color simultaneously. This leads to a total of 15 attribute states and 243 photo configurations.

The graphical interface is a hierarchical menu, where each attribute is a top-level menu entry, and each attribute state is in the corresponding submenu. By clicking a top-level menu, the submenu expands and thus becomes visible and selectable. Only one menu can be expanded at any given time.

The photo attribute states correspond to the current input state and the target state. The interface agent selects an attribute menu to open. Its goal is to reduce the number of clicks necessary to change an attribute, e.g., from two user interactions (filter->color) to one (color). For the *user agent*, the higher level selects a target slot, and the lower level moves to the corresponding location.

Qualitative Policy Inspection

We observe that the interface agent intelligently decides which submenu to open next. We notice that this is never a menu the user recently interacted with as the probability of having to change, e.g., the color twice in a row is minimal and only a result of errors.

9.7 DISCUSSION

MARLUI models the interaction with point-and-click adaptive interfaces as a multi-agent cooperative game by teaching a simulated user agent and an interface agent to cooperate. Learned policies of the interface agent have shown their capability to effectively assist real users. Demonstrating our approach in a wide variety of use cases is a first step towards general methods that are not tied to specific applications nor dependent on manually crafted rules or offline user data collection. However, there are limitations that require further research.

In this work, we have focused on point-and-click interfaces. However, it would be interesting to extend the user agent to model other interaction paradigms. By enhancing our user agent to replicate behavior for other interaction types than clicking, we could extend the possible use cases that MARLUI can support. For example, research has shown that gaze-based selection, similar to cursor movement, follows Fitts' Law [229]. Furthermore, similar concepts can also be applied to human-robot interactions, such as using simulated humans to train human-robot handshakes or human-to-robot handovers [46, 48, 49].

Moreover, user goals can change during human-computer interaction, particularly in creative tasks where users constantly adjust their objective based on intermediate results. This presents challenges for standard RL approaches, which assume goals to remain stationary. Future research on MARL for AUIs needs to focus on finding strategies to easily adapt trained interfaces to changing or new user goals. This is required to establish more robust and flexible adaptive interfaces that can support real-world use cases.

We have demonstrated that our formulation solves problems with up to 5 billion possible states, as in the character creation application (Sec. 9.5.1). However, the complexity of the problem grows exponentially with the number of states. This makes it challenging for MARLUI to scale to interfaces with even larger state spaces. To overcome this, we could explore different input modalities, such as representing the state of the UI as an image instead of using one-hot encoding. This approach is similar to work on RL agents playing video games [183], which showed that image representations can effectively cope with large state spaces.

We have shown that the simulated user agent's behavior was sufficiently human-like to enable the interface agent to learn helpful policies that transfer to real users. The interface agent's performance is inherently limited by the user agent. Therefore, increasing realism in the model of the simulated user is an interesting future research direction, for instance, modeling human-like search [42] or motor control with a biomechanical model [73]. Along similar lines, our work helps with the creation of policies that adapt interfaces given user interactions. However, it does not adapt to

the user themselves (e.g., different levels of expertise). Such personalization is an interesting direction for future research.

Our framework has theoretical appeal because it provides a plausible model of the bilateral nature of AUIS: the adaptation depends on the user, whereas also the user action depends on the adaptation. Modeling this unilaterally as in supervised learning does not reflect reality well. Related to ongoing research [191], we believe that future work can leverage our framework to gain a better theoretical understanding of how users interact with a UI. Our setup has the potential to scale to multiple users with different skills and intentions. This could lead to bespoke assistive UIs for users with specific needs or UIs for users with specific expertise levels.

Finally, our proposed framework enables adaptive policies for different point-and-click tasks and interfaces. We have shown how our framework produces policies that support users in a variety of these interfaces and tasks. Building on this, future work can investigate the transition from framework to developer tool. Tools that enable developers to use our framework easily and consistently will streamline the development process of AUIS.

The question addressed in this paper is whether it is possible to develop a *general framework* for point-and-click AUIS that does not depend on task-specific heuristics or data to generate policies offline. To this end, we have introduced *MARLUI*, a multi-agent reinforcement learning approach. Our method features a user agent and an interface agent. The user agent aims to achieve a task-dependent goal as quickly as possible, while the interface agent learns the underlying task structure by observing the interactions between the user agent and the UI. Since the user agent is RL-based and thus learns through trial-and-error interactions with the interface, it does not require real user data. We have evaluated our approach in simulation and with participants, by replacing the user agent with real users, across five different interfaces and various underlying task structures. The tasks ranged from assigning items to a toolbar, handing out-of-reach objects to the user, selecting the best-performing interface, providing the correct object to the user, and enabling more efficient interaction with a hierarchical menu. Results show that our framework enables the development of AUIS with minimal adjustments while being able to assist real users in their task. We believe that *MARLUI*, and a multi-agent perspective in general, is a promising step towards tools for developing adaptive interfaces, thereby reducing the overhead of developing adaptive strategies on an interface- and task-specific basis.

SUMMARY & INSIGHTS

SUMMARY In summary, we introduced two optimal control strategies: *MagPen* and *MARLUI*. In "Optimal Control for Electromagnetic Haptic Guidance Systems" (Chapter 8), we developed an optimization method for electromagnetic haptic guidance systems using the MPCC framework. This method balanced user input and system control, allowing for trajectory and speed adjustments as needed. It optimized system states and inputs over a receding horizon by solving a stochastic optimal control problem at each timestep. Our design dynamically adjusted forces and automatically modified magnet position and strength, suitable for real-time optimization. We implemented our method on a prototype hardware platform, *MagPen*, and demonstrated that users responded well to the feedback, achieving higher accuracy compared to open-loop and time-dependent closed-loop methods. This approach offered a principled method for haptic guidance, broadly applicable in areas such as drawing, sketching, writing, or virtual haptic tools. However, our MPCC-based framework relied on known system dynamics, which are not always available, especially when considering user behavior.

In our second project, "Multi-Agent Reinforcement Learning for Point-and-Click Adaptive User Interfaces" (Chapter 9), we addressed this limitation by introducing a model-free multi-agent reinforcement learning strategy for adaptive user interfaces. *MARLUI* incorporated a user agent and an interface agent. The user agent aimed to achieve a task-specific goal quickly, while the interface agent learned the task structure by observing the user agent's interactions with the UI. This method did not require real user data, as the user agent learned through trial and error. We evaluated our approach through simulations and real user tests across five different interfaces and various task structures. The results indicated that our framework supported adaptive UI development with minimal adjustments, effectively assisting real users in their tasks. *MARLUI*, and the multi-agent approach in general, represent a promising direction for developing adaptive interfaces, reducing the need for interface- and task-specific strategies.

Overall, our methods enabled an autonomy-automation trade-off, resulting in better user performance in terms of speed, accuracy, and number of actions compared to alternatives. Our research is a starting point for considering and investigating the role of predictive user models in control loops.

IMPLICATIONS Previous systems commonly employ open-loop control [278] to manage intelligent systems. This method does not take the user into account, thereby preventing the balancing of user agency with system automation. Alternatively, some systems rely on heuristics [33, 165, 250, 254], which involve the creation of rules by experts. Additionally, supervised learning [27, 69, 146, 173, 178, 181, 208, 238, 239] is widely used for controlling intelligent systems. However, supervised learning focuses on short-term decisions. By only optimizing for the immediate next step, they limit the system’s ability to intelligently balance user agency with system automation.

In contrast, our approach minimizes a cost function over a receding horizon constrained by user and system dynamics. We have demonstrated how to integrate models of human behavior explicitly and implicitly into optimal control strategies for intelligent systems. This approach allows systems to consider future states and actions and optimize inputs accordingly. In *MagPen*, we embedded a user model explicitly in the control strategy, considering user behavior over a horizon, though mathematically formulating this behavior can be challenging. Conversely, in *MARLUI*, we took a multi-agent RL approach where user behavior is implicitly learned, eliminating the need for explicit behavioral models or system dynamics descriptions, and allowing the method to generalize across interfaces and tasks.

Our method builds on existing cognitive models in the literature, leveraging previous research. This grounding in cognitive models enables advancements on two fronts: 1) developing cognitive models for user bounds, and 2) applying RL methods to more complex scenarios. Treating HCI, and specifically the interaction with intelligent systems, as a multi-agent reinforcement problem is intuitive and facilitates future extensions of our work.

LIMITATIONS Our research has limitations in the context of shared variable interfaces.

First, both *MagPen* and *MARLUI* require an objective function. In task-driven scenarios, these objective functions are relatively intuitive, such as accuracy or task completion time. However, in open-ended scenarios, such as browsing the web or engaging in creative tasks, defining an appropriate objective is challenging. Further research, such as Inverse Reinforcement Learning, is required to extract objective functions from user data.

Second, our work does not explicitly account for how human behavior changes based on the system’s actions. In *MagPen*, we assume a compliant user, and in *MARLUI*, the interaction might be implicitly learned. However, explicit models focusing on the interaction and the influence of system actions on user actions are vital.

Finally, neither *MagPen* nor *MARLUI* considers the user's learning phase. During interaction, users develop a mental model of the system. Ideally, the system would be aware of the user's mental model. Embedding theory-of-mind models into a control strategy could create methods that work for both novice and expert users and personalize interactions based on user behavior, resulting in more intuitive and efficient interactions.

Part IV

CONCLUSION

CONCLUSION

In this dissertation, we investigate the trade-off between user agency and system automation in interactive intelligent systems, focusing specifically on interactions involving shared variables. These interactions allow both a user and an artificial agent to manipulate the same variable. Despite rapid advancements in contextual intelligent systems, research on their interaction remains insufficient. We approach this problem from two angles: first, by investigating a haptic interface as the interaction medium, and second, by leveraging optimal control strategies to balance agency and automation more effectively.

THE DESIGN OF SHARED VARIABLE INTERFACES

In a shared variable interface, both the user and the system act on a shared variable, which can take forms such as kinesthetic haptic devices or graphical menus. Kinesthetic haptic feedback serves as a special instance of this shared variable, where both action and perception occur in the same modality from the user's perspective, such as using hands to change or perceive a joystick's position. This contrasts with interfaces where perception and interaction modalities differ, like touch interfaces. Previous kinesthetic haptic feedback systems often involved complex devices requiring user instrumentation, which limited full user agency. To overcome these limitations, we introduce novel haptic interfaces that elicit grounded forces for system automation and are untethered for full user agency.

1. Contact-free Non-Planar Haptics with a Spherical Electromagnet

In Chapter 5, we introduce *Omni-v1*, a contact-free, non-planar haptic feedback device with a spherical electromagnetic actuator. The actuator consists of three orthogonal coils integrated into a single sphere, capable of eliciting radial and tangential forces on a permanent magnet embedded in a tool. Despite the tool being untethered, the force remains grounded via the electromagnet. Our device can produce forces up to 1N, which can be increased by enhancing the volume of the embedded permanent magnet, albeit at a weight trade-off. Users were able to discriminate at least 25 points spread evenly on the device's sphere (Figure 5.8), with mean errors of $2.5\text{mm} \pm 1.4$, $5.7\text{mm} \pm 4.6$, $6.5\text{mm} \pm 5.2$, and $7.2\text{mm} \pm 5.1$ across

different azimuth angles (Figure 5.9). *Omni-v1* showed that enabling feedback in the same modality as the user’s exertion and with a shared variable results in natural interactions. However, the reliance on external tracking for tool position remained a limitation due to its expense and complexity.

2. Volumetric Sensing and Actuation of Passive Magnetic Tools for Dynamic Haptic Feedback

Chapter 6 introduces on-device sensing with *Omni-v2*, integrating Hall sensors into the *Omni-v1* base. This core contribution allows us to decompose the natural interference caused by simultaneous magnetic tracking and actuation. Our novel gradient-based optimization method minimizes the difference between estimated and observed magnetic fields, affording 3D tracking capabilities with a mean error of 6.9mm during actuation (Figure 6.5). Furthermore, we improved the actuator design compared to *Omni-v1*, doubling the peak force to $\pm 2N$ at 15A (Figure 6.7). We demonstrated *Omni-v2*’s capabilities through three Virtual Reality applications: sculpting, gaming, and soft object exploration (Figure 6.10). *Omni-v2* showed how it is possible to sense the tool state and take it into account in the control loop without user instrumentation or cumbersome setups.

Implications

Our work lays the foundation for investigating haptic devices from a new perspective, that of shared variables interfaces. We have taken the first step towards understanding haptic shared variables interfaces and demonstrate new applications in gaming, object exploration, and design. This integration of sensing and actuating the tool through electromagnetism is both cost-effective and efficient. The primary advantages of *Omni-v1* and *Omni-v2* are their high accuracy and substantial force output, achieved without mechanically moving parts, thereby avoiding wear and increasing user agency. By eliciting forces on the tool manipulated by the user, we enable haptic feedback in the same modality as the user’s exertion, resulting in natural interactions. However, more intelligent system control is needed for complex dynamics.

THE CONTROL OF SHARED VARIABLE INTERFACES

The intelligent control of interfaces with shared variables is crucial to trade off user agency with system automation. Previous systems typically used methods that either do not take the user into account at all or do not predict user behavior. Therefore, previous methods have no or limited ability to intelligently trade off user agency

and system automation. In contrast, we embed, explicitly and implicitly, a predictive model of user behavior into our control strategy. A predictive model allows the system to take future human decisions into account in the control loop. Using future predictions enables a better balance between user agency and system automation.

3. *Optimal Control for Electromagnetic Haptic Guidance Systems*

In Chapter 8, we explore Model Predictive Control to enable haptic guidance for electromagnetic systems. Our real-time approach assists users in tasks like calligraphy (Figure 8.9), sketching (Figure 8.10), or designing (Figure 8.11) by iteratively predicting the motion of an input device (such as a pen) and adjusting the position and strength of an underlying dynamic electromagnetic actuator. In addition to the control strategy, we introduced a prototypical haptic feedback device that allowed us to evaluate our approach (Figure 8.2). Our user study showed that users were more accurate with our method ($6.2\text{mm} \pm 0.8$) compared to open-loop ($38.0\text{mm} \pm 56.9$) or time-dependent closed-loop methods ($45.0\text{mm} \pm 50.8$), as illustrated in Figure 8.8 and Table 8.4. *MagPen* demonstrates that an explicit predictive user model and system dynamics with a receding horizon cost function enable the control of a haptic system that trades off user autonomy with system automation. However, the reliance on known system dynamics posed a challenge, especially in embedding user behavior.

4. *Multi-Agent Reinforcement Learning for Point-and-Click Adaptive User Interfaces*

In Chapter 9, we introduce a model-free reinforcement learning approach to overcome the need for explicit system dynamics. We also proposed adaptive user interfaces that automatically adjust based on user context and previous inputs. In our MARL formulation, a user agent simulates real user interactions, while an interface agent learns to optimize the UI based on the user agent's performance. This approach significantly reduced the number of actions users needed (3.34) compared to supervised-learning (3.87) and static baselines (5.73), as illustrated in Figure 9.7 and Table 9.1. Unique to our approach is that it requires minimal tuning between tasks and scenarios. To show this, we introduced four additional use cases: Interface Selection (Figure 9.8), Block Building (Figure 9.9), Out-of-reach Item Grabbing (Figure 9.10), and Photo Editing (Figure 9.11). *MARLUI* demonstrates that it is possible to learn implicit models of human behavior and tasks by treating interaction as a multi-agent game. These implicit models can then be integrated into intelligent control strategies that take human decision-making into account.

Implications

We demonstrate how to integrate models of human behavior explicitly and implicitly into optimal control strategies for intelligent systems. This approach allows systems to consider future states and actions and optimize inputs accordingly. In *MagPen*, we embed a user model explicitly in the control strategy, considering user behavior over a horizon, though mathematically formulating this behavior can be challenging. Conversely, in *MARLUI*, we take a multi-agent RL approach where user behavior is implicitly learned, eliminating the need for explicit behavioral models or system dynamics descriptions, and allowing the method to generalize across interfaces and tasks. Both approaches successfully balance user agency and automation, assisting users in their tasks.

CONCLUSION

We started with the question: “*How can we algorithmically control intelligent systems with shared variables to balance user agency and system automation?*” To answer this question, we provide four distinct contributions in two parts: the design and control of shared variable interfaces. We demonstrate that integrating models of human behavior into control strategies, either explicitly or implicitly, enables intelligent systems to take user agency better into account. Furthermore, we show that user agency versus system automation is not just an algorithmic problem, and that this trade-off must also be taken into account during the engineering of physical devices and the design of interfaces. Hence, we advocate for an integrated end-to-end approach for the interaction with intelligent systems that takes an algorithmic, engineering, and design perspective.

OUTLOOK AND FUTURE WORK

This dissertation has demonstrated the potential of shared control for interaction with intelligent systems that have shared variables, but further exploration is needed. We identify three important directions for future research: improving the system’s understanding of human behavior, enabling systems to reason from a human perspective, and enhancing control strategies to incorporate models of human behavior.

12.1 UNDERSTANDING HUMANS

12.1.1 *Human Sensing & Inference*

The shift towards contextual intelligent systems will integrate the world at large—including the users themselves—into the interface. For instance, in a scenario where a robot and a human collaborate to clean an apartment, the robot should infer and predict the user’s intent from the state of the world and the user’s actions without explicit instructions.

Accurate human state estimation is essential for this, including physical states captured through computer vision and, more challengingly, latent states such as expertise, fatigue, and intent. Advances in machine learning, particularly in areas such as human latent state inference [153], are needed to tackle these challenges. Additionally, the system should be aware of its environment and capable of extrapolating and predicting future states. Physics-based machine learning holds promise for understanding physical world dynamics [49], while behavioral models are necessary for predicting human behavior. A first project could be to learn a classifier that predicts the extent to which a user is an expert in complex software such as CAD based on their cursor movement.

12.1.2 *Human Behavioral Models*

Sensing the human state alone is insufficient. To apply current state estimations to intelligent control, we need to predict future human states. This requires behavior models that capture the complex dynamics of human behavior. Challenges include stochastic system dynamics, changing user optimization goals, and the impact of system actions on users.

Data-driven approaches such as imitation learning and inverse reinforcement learning, combined with cognitive models, can help tackle these problems [138]. Reinforcement learning, assuming computational rationality, can approximate optimal human-like behavior. By incorporating data-driven priors, we can model more complex behaviors [8]. Integrating these approaches will enhance our understanding of human behavior and improve interface design and interaction paradigms. An interesting direction could be to use text-conditioned models that generate realistic user behavior given a prompt (e.g., "A novice user browses the website to look for new sneakers").

12.1.3 *Non-Stationary and Abstract Goals*

Our research has assumed specific, concrete goals, but this may not translate well to more abstract tasks such as creative activities or web browsing. Goals can also change dynamically during an activity.

Future research should focus on hierarchically structuring goals and dynamically adjusting their probabilities. Recognizing interactions as stochastic decisions at various levels, a hierarchical goal model could better adapt to uncertainty. Additionally, understanding the necessary accuracy of higher-level goals is crucial. In *MagPen*, we showed that a low-level polynomial approximation of a sketch was sufficient when updated iteratively.

12.2 UNDERSTANDING WHAT HUMANS UNDERSTAND

The systems and humans might observe the world differently. Whether this is due to occlusions, quality of sensors, or other reasons, the system should be able to infer i) what the user has observed and ii) what the user believes about the world state. These capabilities might allow the system to predict future user actions and expectations more accurately.

12.2.1 *Human-Like Scene Understanding*

Humans do not have perfect scene understanding. They might miss something outside their field of view or something that does not match their prior knowledge. Current user models often assume complete scene awareness, which can be problematic for human-like user models. Future studies should integrate insights from human cognitive perception with advanced computer vision techniques. Research questions include determining what a person can see from a monocular RGB image

or video and estimating their field of view given a map and location. Extending gaze estimation techniques to real-world scenarios could provide valuable insights [6].

12.2.2 *Theory-of-Mind User Models*

Understanding user beliefs is crucial, but integrating these beliefs into control systems remains a challenge. Future research could explore incorporating probabilistic belief inferences into control frameworks. This requires precise belief inference, representation, and dynamics, along with optimization criteria that meaningfully utilize these beliefs.

Bayesian models offer a tractable path, keeping track of belief states at each inference step [15]. However, this requires knowing probability matrices. Recurrent networks such as LSTM can approximate complex dynamics, but the exact solution to the belief model problem remains an open question. A first project could predict human-belief states in simple games in which the user cooperates with an agent to solve a task (e.g., finding the exit of a maze), where the belief is limited to the agent state (e.g., position) and end-goal (the maze exit).

12.3 COOPERATIVE CONTROL

Finally, we need to integrate human sensing, inference, and behavioral models into intelligent control strategies, paving the way for more intuitive and seamless interactions between humans and machines that respect and enhance human decision-making processes.

12.3.1 *Embedding User Models in System Control*

Core to useful cooperative control strategies is embedding models of human behavior into the system control. However, how to embed human behavioral models is not entirely clear. We could embed these human models hierarchically into the system control strategy. However, this would be computationally expensive as we need to solve multiple control problems per timestep. Alternatively, we can learn implicit user models either from data or in simulation. While this might be a more scalable approach, how to learn these models implicitly remains an open question. A first project would be to directly compare explicit versus implicit embedded user models in a simple task to investigate the trade-offs involved.

12.3.2 *Combine Model-based and Model-free approaches*

Model-based methods like MPC and model-free methods like RL each have unique advantages. MPC can adjust to new and different objectives without retraining, while RL excels at developing policies for unpredictable system behaviors. Combining these approaches could greatly improve computational assistance for users.

Future developments might involve using MPC to offer initial direction in novel situations based on user goals, with RL continuously refining and customizing the system. Learning complex environmental dynamics and integrating them into model-predictive optimization can compute policies without retraining for new objectives. Alternatively, inverse reinforcement learning might offer cost functions to use in MPC when system dynamics are known.

BIBLIOGRAPHY

1. Abbink, D. A., Mulder, M. & Boer, E. R. Haptic shared control: smoothly shifting control authority? *Cognition, Technology & Work* **14**, 19. doi:10.1007/s10111-011-0192-5 (2012).
2. Abler, A., Zárata, J., Langerak, T., Vechev, V. & Hilliges, O. *Hedgehog: Handheld Spherical Pin Array based on a Central Electromagnetic Actuator* in *2021 IEEE World Haptics Conference (WHC)* (2021), 133. doi:10.1109/WHC49131.2021.9517197.
3. Abut, T. & Soygüder, S. Interface design and performance analysis for a haptic robot. *Muş Alparslan Üniversitesi Fen Bilimleri Dergisi* **6**, 553 (2018).
4. Adel, A., Mansour, M., Micheal, M. M., Abdelmawla, A., Khalil, I. S. M. & Misra, S. Magnetic Localization for an Electromagnetic-Based Haptic Interface. *IEEE Magnetics Letters* **10**, 1. doi:10.1109/LMAG.2019.2908149 (2019).
5. Aguiar, A. P., Hespanha, J. P. & Kokotovic, P. V. Performance limitations in reference tracking and path following for nonlinear systems. *Automatica* **44**, 598. doi:10.1016/j.automatica.2007.06.030 (2008).
6. Akinyelu, A. A. & Blignaut, P. Convolutional neural network-based methods for eye gaze estimation: A survey. *IEEE Access* **8**, 142581. doi:10.1109/access.2020.3013540 (2020).
7. Aksan, E., Pece, F. & Hilliges, O. *DeepWriting: Making Digital Ink Editable via Deep Generative Modeling* in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (eds Mandryk, R. L., Hancock, M., Perry, M. & Cox, A. L.) (Association for Computing Machinery, Montreal QC, Canada, 2018), 1. doi:10.1145/3173574.3173779.
8. Albaba, M., Christen, S., Gebhardt, C., Langarek, T., Black, M. J. & Hilliges, O. *RILE: Reinforced Imitation Learning* 2024.
9. Araujo, B., Jota, R., Perumal, V., Yao, J. X., Singh, K. & Wigdor, D. *Snake Charmer: Physically Enabling Virtual Objects* in *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction* (eds Bakker, S., Hummels, C., Ullmer, B., Geurts, L., Hengeveld, B., Saakes, D. & Broekhuijsen, M.) (ACM, Eindhoven, Netherlands, 2016), 218. doi:10.1145/2839462.2839484.

10. Ashbrook, D., Baudisch, P. & White, S. *Nenya: Subtle and Eyes-Free Mobile Input with a Magnetically-Tracked Finger Ring* in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (eds Tan, D. S., Amershi, S., Begole, B., Kellogg, W. A. & Tungare, M.) (Association for Computing Machinery, Vancouver, BC, Canada, 2011), 2043. doi:10.1145/1978942.1979238.
11. Åström, K. J. Optimal control of Markov processes with incomplete state information. *Journal of mathematical analysis and applications* **10**, 174. doi:10.1016/0022-247x(65)90154-x (1965).
12. Åström, K. & Häggglund, T. *Advanced PID Control* (ISA - The Instrumentation, Systems and Automation Society, 2006).
13. Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., Bonnefon, J.-F. & Rahwan, I. The moral machine experiment. *Nature* **563**, 59. doi:10.1038/s41586-018-0637-6 (2018).
14. Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B. & Mordatch, I. Emergent tool use from multi-agent autotutorials. *arXiv preprint arXiv:1909.07528*. doi:10.48550/arxiv.1909.07528 (2019).
15. Baker, C., Saxe, R. & Tenenbaum, J. *Bayesian theory of mind: Modeling joint belief-desire attribution* in *Proceedings of the annual meeting of the cognitive science society* (eds Carlson, L. A., Hölscher, C. & Shipley, T. F.) **33** (Wiley, 2011).
16. Bansal, G., Nushi, B., Kamar, E., Lasecki, W. S., Weld, D. S. & Horvitz, E. *Beyond accuracy: The role of mental models in human-AI team performance* in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* **7** (Association for the Advancement of Artificial Intelligence (AAAI), 2019), 2. doi:10.1609/hcomp.v7i1.5285.
17. Bansal, G., Nushi, B., Kamar, E., Weld, D. S., Lasecki, W. S. & Horvitz, E. *Updates in human-AI teams: Understanding and addressing the performance/compatibility tradeoff* in *Proceedings of the AAAI Conference on Artificial Intelligence* **33** (Association for the Advancement of Artificial Intelligence, 2019), 2429. doi:10.1609/aaai.v33i01.33012429.
18. Bar-Hillel, Y. The present status of automatic translation of languages. *Advances in computers* **1**, 91. doi:10.1016/s0065-2458(08)60607-5 (1960).

19. Benko, H., Holz, C., Sinclair, M. & Ofek, E. *NormalTouch and TextureTouch: High-fidelity 3D Haptic Shape Rendering on Handheld Virtual Reality Controllers* in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (ACM, Tokyo, Japan, 2016), 717. doi:10.1145/2984511.2984526.
20. Benner, P. *Control theory basics* 2007.
21. Berkelman, P., Bozlee, S. & Miyasaka, M. Interactive rigid-body dynamics and deformable surface simulations with co-located maglev haptic and 3D graphic display. *International journal on advances in intelligent systems* **6**, 289 (2013).
22. Berkelman, P. & Dzadovsky, M. *Extending the motion ranges of magnetic levitation for haptic interaction* in *World Haptics 2009-Third Joint EuroHaptics conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems* (IEEE, 2009), 517. doi:10.1109/whc.2009.4810897.
23. Berkelman, P., Miyasaka, M. & Anderson, J. *Co-located 3D graphic and haptic display using electromagnetic levitation* in *2012 IEEE Haptics Symposium (HAPTICS)* (IEEE, 2012), 77. doi:10.1109/haptic.2012.6183773.
24. Berkelman, P., Tix, B. & Abdul-Ghani, H. Electromagnetic position sensing and force feedback for a magnetic stylus with an interactive display. *IEEE Magnetics Letters* **10**, 1. doi:10.1109/lmag.2018.2886339 (2018).
25. Berkelman, P. J., Butler, Z. J. & Hollis, R. L. *Design of a hemispherical magnetic levitation haptic interface device* in *Proceedings of the ASME Winter Annual Meeting, Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, (Atlanta)* (American Society of Mechanical Engineers, 1996). doi:10.1115/imece1996-0373.
26. Berkelman, P. J. & Hollis, R. L. Lorentz magnetic levitation for haptic interaction: Device design, performance, and integration with physical simulations. *The International Journal of Robotics Research* **19**, 644. doi:10.1177/027836490001900703 (2000).
27. Berry, P. M., Gervasio, M., Peintner, B. & Yorke-Smith, N. PTIME: Personalized assistance for calendaring. *ACM Trans. Intell. Syst. Technol.* **2**. doi:10.1145/1989734.1989744 (2011).
28. Bhadra, N., Peckham, P. H., Keith, M. W., Kilgore, K. L., Montague, F., Gazdik, M. & Stage, T. Implementation of an implantable joint-angle transducer. *Journal of rehabilitation research and development* **39**, 411 (2002).

29. Bosma, W. & André, E. *Exploiting Emotions to Disambiguate Dialogue Acts in Proceedings of the 9th International Conference on Intelligent User Interfaces* (eds Vanderdonckt, J., Nunes, N. J. & Rich, C.) (Association for Computing Machinery, Funchal, Madeira, Portugal, 2004), 85. doi:10.1145/964442.964459.
30. Botvinick, M. M. Hierarchical reinforcement learning and decision making. *Current opinion in neurobiology* **22**, 956. doi:10.1016/j.conb.2012.05.008 (2012).
31. Brink, J. B., Petruska, A. J., Johnson, D. E. & Abbott, J. J. *Factors affecting the design of untethered magnetic haptic interfaces in 2014 IEEE Haptics Symposium (HAPTICS)* (IEEE, 2014), 107. doi:10.1109/haptics.2014.6775441.
32. Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J. & Zaremba, W. *OpenAI Gym* 2016.
33. Browne, D., Totterdell, P. & Norman, M. *Adaptive user interfaces* (Elsevier, 2016).
34. Camacho, E. F. & Alba, C. B. *Model predictive control* (Springer Science & Business Media, 2013).
35. Card, S., Moran, T. & Newell, A. *The Psychology of Human Computer Interaction* 1983.
36. Card, S. K. The Model Human Processor: A Model for Making Engineering Calculations of Human Performance. *Proceedings of the Human Factors Society Annual Meeting* **25**, 301. doi:10.1177/107118138102500180 (1981).
37. Card, S. K., Moran, T. P. & Newell, A. The Keystroke-Level Model for User Performance Time with Interactive Systems. *Commun. ACM* **23**, 396. doi:10.1145/358886.358895 (1980).
38. Cheema, N., Frey-Law, L. A., Naderi, K., Lehtinen, J., Slusallek, P. & Hämmäläinen, P. *Predicting mid-air interaction movements and fatigue using deep reinforcement learning in Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (ACM, 2020), 1. doi:10.1145/3313831.3376701.
39. Chen, G. A New Framework for Multi-Agent Reinforcement Learning–Centralized Training and Exploration with Decentralized Execution via Policy Distillation. *arXiv preprint arXiv:1910.09152*. doi:10.48550/arxiv.1910.09152 (2019).

40. Chen, J. Y. C. & Barnes, M. J. Human–Agent Teaming for Multirobot Control: A Review of Human Factors Issues. *IEEE Transactions on Human-Machine Systems* **44**, 13. doi:10.1109/THMS.2013.2293535 (2014).
41. Chen, M., Beutel, A., Covington, P., Jain, S., Belletti, F. & Chi, E. H. *Top-K Off-Policy Correction for a REINFORCE Recommender System* in *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining* (eds Culpepper, J. S., Moffat, A., Bennett, P. N. & Lerman, K.) (ACM, 2019), 456. doi:10.1145/3289600.3290999.
42. Chen, X., Bailly, G., Brumby, D. P., Oulasvirta, A. & Howes, A. *The emergence of interactive behavior: A model of rational menu search* in *Proceedings of the 33rd annual ACM conference on human factors in computing systems* (eds Begole, B., Kim, J., Inkpen, K. & Woo, W.) (Centre National de la Recherche Scientifique, 2015), 4217. doi:10.1145/2702123.2702483.
43. Chen, K.-Y., Patel, S. N. & Keller, S. *Finexus: Tracking Precise Motions of Multiple Fingertips Using Magnetic Sensing* in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (eds Kaye, J., Druin, A., Lampe, C., Morris, D. & Hourcade, J. P.) (Association for Computing Machinery, San Jose, California, USA, 2016), 1504. doi:10.1145/2858036.2858125.
44. Choi, I. & Follmer, S. *Wolverine: A Wearable Haptic Interface for Grasping in VR* in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (ACM, Tokyo, Japan, 2016), 117. doi:10.1145/2984751.2985725.
45. Choi, I., Ofek, E., Benko, H., Sinclair, M. & Holz, C. *CLAW: A Multifunctional Handheld Haptic Controller for Grasping, Touching, and Triggering in Virtual Reality* in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (eds Mandryk, R. L., Hancock, M., Perry, M. & Cox, A. L.) (Association for Computing Machinery, Montreal QC, Canada, 2018), 1. doi:10.1145/3173574.3174228.
46. Christen, S., Feng, L., Yang, W., Chao, Y.-W., Hilliges, O. & Song, J. SynH2R: Synthesizing Hand-Object Motions for Learning Human-to-Robot Handovers. *arXiv preprint arXiv:2311.05599*, 3168. doi:10.48550/arxiv.2311.05599 (2023).
47. Christen, S., Jendele, L., Aksan, E. & Hilliges, O. Learning Functionally Decomposed Hierarchies for Continuous Control Tasks With Path Planning. *IEEE Robotics and Automation Letters* **6**, 3623. doi:10.1109/LRA.2021.3060403 (2021).

48. Christen, S., Stevšić, S. & Hilliges, O. *Demonstration-Guided Deep Reinforcement Learning of Control Policies for Dexterous Human-Robot Interaction* in *2019 International Conference on Robotics and Automation (ICRA)* (2019), 2161. doi:10.1109/ICRA.2019.8794065.
49. Christen, S., Yang, W., Pérez-D'Arpino, C., Hilliges, O., Fox, D. & Chao, Y.-W. *Learning human-to-robot handovers from point clouds* in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, 2023), 9654. doi:10.1109/cvpr52729.2023.00931.
50. Cockburn, A., Gutwin, C. & Greenberg, S. *A Predictive Model of Menu Performance* in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Association for Computing Machinery, San Jose, California, USA, 2007), 627. doi:10.1145/1240624.1240723.
51. COMSOL Multiphysics. *COMSOL Multiphysics* www.comsol.com.
52. Costello, R. G. The surge model of the well-trained human operator in simple manual control. *IEEE Transactions on Man-Machine Systems* **9**, 2. doi:10.1109/tmms.1968.300028 (1968).
53. Craik, K. J. W. Theory of the human operator in control systems. *British Journal of Psychology. General Section* **38**, 56. doi:https://doi.org/10.1111/j.2044-8295.1947.tb01141.x (1947).
54. Culbertson, H., Schorr, S. B. & Okamura, A. M. Haptics: The Present and Future of Artificial Touch Sensation. *Annual Review of Control, Robotics, and Autonomous Systems* **1**, 385. doi:10.1146/annurev-control-060117-105043 (2018).
55. CyberGlove Systems Inc. *CyberGrasp Glove* <http://www.cyberglovesystems.com/cybergrasp>. Last accessed: 12.03.2017.
56. CyberGlove Systems Inc. *CyberTouch Glove* <http://www.cyberglovesystems.com/cybertouch>. Last accessed: 18.03.2017.
57. Da Silva, M., Abe, Y. & Popović, J. Simulation of Human Motion Data using Short-Horizon Model-Predictive Control. *Computer Graphics Forum* **27**, 371. doi:10.1111/j.1467-8659.2008.01134.x (2008).
58. Dai, P., Lin, C., Mausam, M. & Weld, D. POMDP-based control of workflows for crowdsourcing. *Artificial Intelligence* **202**, 52. doi:10.1016/j.artint.2013.06.002 (2013).

59. Debard, Q., Dibangoye, J. S., Canu, S. & Wolf, C. *Learning 3D Navigation Protocols on Touch Interfaces with Cooperative Multi-agent Reinforcement Learning in Machine Learning and Knowledge Discovery in Databases* (eds Brefeld, U., Fromont, E., Hotho, A., Knobbe, A., Maathuis, M. & Robardet, C.) (Springer International Publishing, Cham, 2020), 35. doi:10.1007/978-3-030-46133-1_3.
60. Demir, M., Cooke, N. J. & Amazeen, P. G. A conceptual model of team dynamical behaviors and performance in human-autonomy teaming. *Cognitive Systems Research* **52**, 497. doi:10.1016/j.cogsys.2018.07.029 (2018).
61. Demir, M., McNeese, N. J. & Cooke, N. J. Team situation awareness within the context of human-autonomy teaming. *Cognitive Systems Research* **46**, 3. doi:10.1016/j.cogsys.2016.11.003 (2017).
62. Denison, S., Bonawitz, E., Gopnik, A. & Griffiths, T. L. Rational variability in children's causal inferences: The sampling hypothesis. *Cognition* **126**, 285. doi:10.1016/j.cognition.2012.10.010 (2013).
63. Domahidi, A. & Jerez, J. *FORCES Pro: code generation for embedded optimization*. <https://www.embotech.com/FORCES-Pro>. 2016.
64. Duchowski, A. T., Krejtz, K., Krejtz, I., Biele, C., Niedzielska, A., Kiefer, P., Raubal, M. & Giannopoulos, I. *The index of pupillary activity: Measuring cognitive load vis-à-vis task difficulty with pupil oscillation in Proceedings of the 2018 CHI conference on human factors in computing systems* (2018), 1. doi:10.1145/3173574.3173856.
65. Engelbart, D. C. & English, W. K. *A research center for augmenting human intellect in Proceedings of the December 9-11, 1968, fall joint computer conference, part I* **33** (ACM Press, 1968), 395. doi:10.1145/1476589.1476645.
66. Engelbart, D. C. in *Ideas That Created the Future: Classic Papers of Computer Science* (The MIT Press, 2021). doi:10.7551/mitpress/12274.003.0024.
67. Esmailie, F., Cavilla, M. S. & Ameen, T. A. A Thermal Transient Model of Heat Transfer Within an Omnimagnet. *ASME International Mechanical Engineering Congress and Exposition Volume 8: Heat Transfer and Thermal Engineering*, V008T10A046. doi:10.1115/IMECE2017-72120 (2017).
68. Ewart J. de Visser, R. P. & Shaw, T. H. From 'automation' to 'autonomy': the importance of trust repair in human-machine interaction. *Ergonomics* **61**. PMID: 29578376, 1409. doi:10.1080/00140139.2018.1457725 (2018).

69. Faulring, A., Myers, B., Mohnkern, K., Schmerl, B., Steinfeld, A., Zimmerman, J., Smailagic, A., Hansen, J. & Siewiorek, D. *Agent-Assisted Task Management That Reduces Email Overload* in *Proceedings of the 15th International Conference on Intelligent User Interfaces* (Association for Computing Machinery, Hong Kong, China, 2010), 61. doi:10.1145/1719970.1719980.
70. Faulwasser, T., Kern, B. & Findeisen, R. *Model predictive path-following for constrained nonlinear systems* in *Proceedings of the 48th IEEE Conference on Decision and Control (CDC)* (IEEE, 2009), 8642. doi:10.1109/CDC.2009.5399744.
71. Fernando, P., Peiris, R. L. & Nanayakkara, S. *I-Draw: Towards a Freehand Drawing Assistant* in *Proceedings of the 26th Australian Computer-Human Interaction Conference on Designing Futures: The Future of Design* (Association for Computing Machinery, Sydney, New South Wales, Australia, 2014), 208. doi:10.1145/2686612.2686644.
72. Ferretti, S., Mirri, S., Prandi, C. & Salomoni, P. *Exploiting reinforcement learning to profile users and personalize web pages* in *2014 IEEE 38th International Computer Software and Applications Conference Workshops* (IEEE, 2014), 252. doi:10.1109/compsacw.2014.45.
73. Fischer, F., Bachinski, M., Klar, M., Fleig, A. & Müller, J. Reinforcement learning control of a biomechanical model of the upper extremity. *Scientific Reports* **11**, 1. doi:10.1038/s41598-021-93760-1 (2021).
74. Fitts, P. M. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of experimental psychology* **47**, 381. doi:10.1037/h0055392 (1954).
75. Flemisch, F., Abbink, D., Itoh, M., Pacaux-Lemoine, M.-P. & Weßel, G. Shared control is the sharp end of cooperation: Towards a common framework of joint action, shared control and human machine cooperation. *IFAC-PapersOnLine* **49**, 72. doi:10.1016/j.ifacol.2016.10.464 (2016).
76. Foerster, J., Assael, I. A., De Freitas, N. & Whiteson, S. Learning to communicate with deep multi-agent reinforcement learning. *Advances in neural information processing systems* **29** (eds Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I. & Garnett, R.) doi:10.48550/arxiv.1605.06676 (2016).
77. Forsyth, B. A. C. & MacLean, K. E. Predictive Haptic Guidance: Intelligent User Assistance for the Control of Dynamic Tasks. *IEEE Transactions on Visualization and Computer Graphics* **12**, 103. doi:10.1109/TVCG.2006.11 (2006).

78. Frank, M. J. & Badre, D. Mechanisms of Hierarchical Reinforcement Learning in Corticostriatal Circuits 1: Computational Analysis. *Cerebral Cortex* **22**, 509. doi:10.1093/cercor/bhr114 (2011).
79. Franklin, G. F., Powell, D. J. & Emami-Naeini, A. *Feedback Control of Dynamic Systems* 4th (Prentice Hall PTR, USA, 2001).
80. Gašić, M. & Young, S. Gaussian processes for POMDP-based dialogue manager optimization. *IEEE Transactions on Audio, Speech and Language Processing* **22**, 28. doi:10.1109/TASL.2013.2282190 (2014).
81. Gebhardt, C. *Optimal Control to Support High-Level User Goals in Human-Computer Interaction* PhD thesis (ETH Zurich, 2020). doi:10.3929/ethz-b-000479155.
82. Gebhardt, C., Hecox, B., van Opheusden, B., Wigdor, D., Hillis, J., Hilliges, O. & Benko, H. *Learning Cooperative Personalized Policies from Gaze Data* in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (ACM, New Orleans, LA, USA, 2019). doi:10.1145/3332165.3347933.
83. Gebhardt, C., Oulasvirta, A. & Hilliges, O. Hierarchical Reinforcement Learning as a Model of Human Task Interleaving. *Computational Brain and Behavior*. doi:10.1007/s42113-020-00093-9 (2021).
84. Gebhardt, C., Stevšić, S. & Hilliges, O. Optimizing for Aesthetically Pleasing Quadrotor Camera Motion. *ACM Trans. Graph.* **37**, 1. doi:10.1145/3197517.3201390 (2018).
85. Gershman, S. J., Horvitz, E. J. & Tenenbaum, J. B. Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science* **349**, 273. doi:10.1126/science.aac6076 (2015).
86. Gershman, S. J., Vul, E. & Tenenbaum, J. B. Multistability and Perceptual Inference. *Neural Computation* **24**, 1. doi:10.1162/NECO_a_00226 (2012).
87. Gibbs, B. P. *Advanced Kalman filtering, least-squares and modeling: a practical handbook* doi:10.1002/9780470890042 (John Wiley & Sons, 2011).
88. Glowacka, D. *Bandit algorithms in recommender systems* in *Proceedings of the 13th ACM Conference on Recommender Systems* (ACM, 2019), 574. doi:10.1145/3298689.3346956.
89. Gmytrasiewicz, P. J. & Doshi, P. *Interactive pomdps: Properties and preliminary results* in *International Conference on Autonomous Agents: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems- 3* (IEEE Computer Society, 2004), 1374. doi:10.1109/aamas.2004.154.

90. Görge, D. Relations between model predictive control and reinforcement learning. *IFAC-PapersOnLine* **50**, 4920. doi:10.1016/j.ifacol.2017.08.747 (2017).
91. Greenberg, A. M. & Marble, J. L. Foundational concepts in person-machine teaming. *Frontiers in Physics* **10**. doi:10.3389/fphy.2022.1080132 (2023).
92. Gu, X., Zhang, Y., Sun, W., Bian, Y., Zhou, D. & Kristensson, P. O. *Dexmo: An Inexpensive and Lightweight Mechanical Exoskeleton for Motion Capture and Force Feedback in VR* in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (ACM, Santa Clara, California, USA, 2016), 1991. doi:10.1145/2858036.2858487.
93. Guiard, Y. & Rioul, O. *A mathematical description of the speed/accuracy trade-off of aimed movement* in *Proceedings of the 2015 British HCI Conference* (eds Lawson, S. & Dickinson, P.) (ACM, 2015), 91. doi:10.1145/2783446.2783574.
94. Han, X., Seki, H., Kamiya, Y. & Hikizu, M. *Wearable handwriting input device using magnetic field* in *SICE Annual Conference 2007* (2007), 365. doi:10.1109/SICE.2007.4421009.
95. Hayward, V., Astley, O. R., Cruz-Hernandez, M., Grant, D. & Robles-De-La-Torre, G. Haptic interfaces and devices. *Sensor Review* **24**, 16. doi:10.1108/02602280410515770 (2004).
96. Hee Lee, G., Faundorfer, F. & Pollefeys, M. *Motion estimation for self-driving cars with a generalized camera* in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, 2013), 2746. doi:10.1109/cvpr.2013.354.
97. Heo, S., Chung, C., Lee, G. & Wigdor, D. *Thor's hammer: An ungrounded force feedback device utilizing propeller-induced propulsive force* in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (eds Mandryk, R. L., Hancock, M., Perry, M. & Cox, A. L.) (ACM, 2018), 525. doi:10.1145/3173574.3174099.
98. Hick, W. E. On the Rate of Gain of Information. *Quarterly Journal of Experimental Psychology* **4**, 11. doi:10.1080/17470215208416600 (1952).
99. Hinchet, R., Vechev, V., Shea, H. & Hilliges, O. *DextrES: Wearable Haptic Feedback for Grasping in VR via a Thin Form-Factor Electrostatic Brake* in *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (Association for Computing Machinery, Berlin, Germany, 2018), 901. doi:10.1145/3242587.3242657.

100. Holkar, K. & Waghmare, L. M. An overview of model predictive control. *International Journal of control and automation* **3**, 47 (2010).
101. Horvitz, E., Breese, J., Heckerman, D., Hovel, D. & Rommelse, K. *The Lumière Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users* in *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence* (eds Cooper, G. F. & Moral, S.) (Morgan Kaufmann Publishers Inc., Madison, Wisconsin, 1998), 256.
102. Hoshi, T., Takahashi, M., Iwamoto, T. & Shinoda, H. Noncontact tactile display based on radiation pressure of airborne ultrasound. *IEEE Transactions on Haptics* **3**, 155. doi:10.1109/toh.2010.4 (2010).
103. Howard, R. A. Dynamic programming and markov processes. (1960).
104. Howes, A., Chen, X., Acharya, A. & Lewis, R. L. Interaction as an emergent property of a Partially Observable Markov Decision Process. *Computational interaction*, 287. doi:10.1093/oso/9780198799603.003.0011 (2018).
105. Hu, Z., Liang, Y., Zhang, J., Li, Z. & Liu, Y. *Inference aided reinforcement learning for incentive mechanism design in crowdsourcing* in *Advances in Neural Information Processing Systems* **31** (Cornell University, 2018), 5508.
106. Ikkala, A., Fischer, F., Klar, M., Bachinski, M., Fleig, A., Howes, A., Hämäläinen, P., Müller, J., Murray-Smith, R. & Oulasvirta, A. Breathing Life Into Biomechanical User Models. *ACM Symposium on User Interface Software and Technology*. doi:10.1145/3526113.3545689 (2022).
107. Imai, S. *Is github copilot a substitute for human pair-programming? an empirical study* in *Proceedings of the ACM/IEEE 44th International Conference on Software Engineering: Companion Proceedings* (IEEE, 2022), 319. doi:10.1109/icse-companion55297.2022.9793778.
108. Iqbal, H. & Yi, B.-J. Design and Experimental Verification of a 3-DOF Spherical Electromagnetic Brake for Haptic Interface. *International Journal of Control, Automation and Systems* **18**, 1. doi:0.1007/s12555-019-0188-0 (2019).
109. Ismail, Z. H. & Sariff, N. in *Applications of Mobile Robots* (ed Hurtado, E. G.) chap. 1 (IntechOpen, Rijeka, 2018). doi:10.5772/intechopen.79337.
110. Israr, A. & Poupyrev, I. *Tactile Brush: Drawing on Skin with a Tactile Grid Display* in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (eds Tan, D. S., Amershi, S., Begole, B., Kellogg, W. A. & Tungare, M.) (Association for Computing Machinery, Vancouver, BC, Canada, 2011), 2019. doi:10.1145/1978942.1979235.

111. Jackson, J. D. *Classical electrodynamics* (John Wiley & Sons, 2007).
112. Jaderberg, M., Czarnnecki, W. M., Dunning, I., Marris, L., Lever, G., Castañeda, A. G., Beattie, C., Rabinowitz, N. C., Morcos, A. S., Ruderman, A., Sonnerat, N., Green, T., Deason, L., Leibo, J. Z., Silver, D., Hassabis, D., Kavukcuoglu, K. & Graepel, T. Human-level performance in 3D multiplayer games with population-based reinforcement learning. *Science* **364**, 859. doi:10.1126/science.aau6249 (2019).
113. Jagacinski, R. J. & Flach, J. M. *Control Theory for Humans: Quantitative approaches to modeling performance* doi:10.1201/9781315144948 (Lawrence Erlbaum, Mahwah, NJ, 2003).
114. Janai, J., Güney, F., Behl, A., Geiger, A., *et al.* Computer vision for autonomous vehicles: Problems, datasets and state of the art. *Foundations and Trends® in Computer Graphics and Vision* **12**, 1. doi:10.1561/06000000079 (2020).
115. Johansen, T. A. Approximate explicit receding horizon control of constrained nonlinear systems. *Automatica* **40**, 293. doi:10.1016/j.automatica.2003.09.021 (2004).
116. John R. Anderson, M. M. & Lebiere, C. ACT-R: A Theory of Higher Level Cognition and Its Relation to Visual Attention. *Human-Computer Interaction* **12**, 439. doi:10.1207/s15327051hci1204_5 (1997).
117. Johnson, M. & Vera, A. No AI is an island: the case for teaming intelligence. *AI Magazine* **40**, 16 (2019).
118. Jokinen, J., Acharya, A., Uzair, M., Jiang, X. & Oulasvirta, A. *Touchscreen Typing as Optimal Supervisory Control* in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)* (eds Kitamura, Y., Quigley, A., Isbister, K., Igarashi, T., Bjørn, P. & Drucker, S. M.) (ACM, 2021). doi:10.1145/3411764.3445483.
119. Jokinen, J. P. P., Kujala, T. & Oulasvirta, A. Multitasking in Driving as Optimal Adaptation Under Uncertainty. *Human Factors* **63**. PMID: 32731763, 1324. doi:10.1177/0018720820927687 (2021).
120. Ju, W., Bonanni, L., Fletcher, R., Hurwitz, R., Judd, T., Post, R., Reynolds, M. & Yoon, J. *Origami Desk: Integrating Technological Innovation and Human-Centric Design* in *Proceedings of the 4th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques* (Association for Computing Machinery, London, England, 2002), 399. doi:10.1145/778712.778770.

121. Kaelbling, L. P., Littman, M. L. & Cassandra, A. R. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* **101**, 99. doi:[https://doi.org/10.1016/S0004-3702\(98\)00023-X](https://doi.org/10.1016/S0004-3702(98)00023-X) (1998).
122. Kangas, I., Schwoerer, M. & Bernardi, L. *Scalable User Interface Optimization Using Combinatorial Bandits* in *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (ACM, 2022), 3375. doi:[10.1145/3477495.3536325](https://doi.org/10.1145/3477495.3536325).
123. Kelley, C. R. *Manual and Automatic Control* (John Wiley and Sons, Inc., New York, NY, 1968).
124. Kemp, C. C., Anderson, C. D., Nguyen, H., Trevor, A. J. & Xu, Z. *A point-and-click interface for the real world: laser designation of objects for mobile manipulation* in *Proceedings of the 3rd ACM/IEEE International Conference on Human Robot Interaction* (Association for Computing Machinery, Amsterdam, The Netherlands, 2008), 241. doi:[10.1145/1349822.1349854](https://doi.org/10.1145/1349822.1349854).
125. Kianzad, S., Huang, Y., Xiao, R. & MacLean, K. E. *Phasking on Paper: Accessing a Continuum of PHysically Assisted SKetchING* in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (eds Bernhaupt, R., Mueller, F. '., Verweij, D., Andres, J., McGrenere, J., Cockburn, A., Avellino, I., Goguey, A., Bjøn, P., Zhao, S., Samson, B. P. & Kocielnik, R.) (Association for Computing Machinery, Honolulu, HI, USA, 2020), 1. doi:[10.1145/3313831.3376134](https://doi.org/10.1145/3313831.3376134).
126. Kianzad, S. & MacLean, K. E. *Harold's purple crayon rendered in haptics: Large-stroke, handheld ballpoint force feedback* in *2018 IEEE Haptics Symposium (HAPTICS)* (eds Kuchenbecker, K. J., Gerling, G. J. & Visell, Y.) (IEEE, 2018), 106. doi:[10.1109/haptics.2018.8357161](https://doi.org/10.1109/haptics.2018.8357161).
127. Kieras, D. E. & Meyer, D. E. An Overview of the EPIC Architecture for Cognition and Performance With Application to Human-Computer Interaction. *Human-Computer Interaction* **12**, 391. doi:[10.1207/s15327051hci1204_4](https://doi.org/10.1207/s15327051hci1204_4) (1997).
128. Kim, H., Kim, M. & Lee, W. *HapThimble: A Wearable Haptic Device Towards Usable Virtual Touch Screen* in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (ACM, Santa Clara, California, USA, 2016), 3694. doi:[10.1145/2858036.2858196](https://doi.org/10.1145/2858036.2858196).
129. Kim, H., Kim, S., Lee, B., Pak, J., Sohn, M., Lee, G. & Lee, W. *Digital Rubbing: Playful and Intuitive Interaction Technique for Transferring a Graphic Image onto Paper with Pen-Based Computing* in *CHI '08 Extended Abstracts on Human Factors in Computing Systems* (eds Czerwinski, M.,

- Lund, A. M. & Tan, D. S.) (Association for Computing Machinery, Florence, Italy, 2008), 2337. doi:10.1145/1358628.1358680.
130. Kiran, B. R., Sobh, I., Talpaert, V., Mannion, P., Sallab, A. A. A., Yogamani, S. & Pérez, P. Deep Reinforcement Learning for Autonomous Driving: A Survey. *IEEE Transactions on Intelligent Transportation Systems* **23**, 4909. doi:10.1109/TITS.2021.3054625 (2022).
 131. Koch, J., Lucero, A., Hegemann, L. & Oulasvirta, A. *May AI? Design Ideation with Cooperative Contextual Bandits in Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Association for Computing Machinery, Glasgow, Scotland Uk, 2019), 1. doi:10.1145/3290605.3300863.
 132. Kolekar, S. V., Sanjeevi, S. G. & Bormane, D. *Learning style recognition using artificial neural network for adaptive user interface in e-learning in 2010 IEEE International conference on computational intelligence and computing research* (IEEE, 2010), 1. doi:10.1109/icccic.2010.5705768.
 133. Kolmogorov, A. Sulla determinazione empirica di una legge di distribuzione. *Inst. Ital. Attuari, Giorn.* **4**, 83 (1933).
 134. Koyama, Y., Sakamoto, D. & Igarashi, T. Crowd-powered parameter analysis for visual design exploration. *Proceedings of the 27th annual ACM symposium on User interface software and technology - UIST '14* (eds Benko, H., Dontcheva, M. & Wigdor, D.) 65. doi:10.1145/2642918.2647386 (2014).
 135. Koyama, Y., Sakamoto, D. & Igarashi, T. SelPh : Progressive Learning and Support of Manual Photo Color Enhancement. *Proc. of CHI '16*. doi:10.1145/2858036.2858111 (2016).
 136. Krstic, M. On the applicability of PID control to nonlinear second-order systems. *National Science Review* **4**, 668. doi:10.1093/nsr/nwx048 (2017).
 137. Kuo, H.-C., Liang, R.-H., Lin, L.-F. & Chen, B.-Y. *GaussMarbles: Spherical Magnetic Tangibles for Interacting with Portable Physical Constraints in Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (Association for Computing Machinery, San Jose, California, USA, 2016), 4228. doi:10.1145/2858036.2858559.
 138. Kwon, M., Daptardar, S., Schrater, P. R. & Pitkow, X. Inverse rational control with partially observable continuous nonlinear dynamics. *Advances in neural information processing systems* **33**, 7898 (2020).

139. Lafuente, M., Elizondo, S., Fernández, U. J. & Marzo, A. *Comparing a Mid-air Two-Hand Pinching Point-and-Click Technique with Mouse, Keyboard and TouchFree* in *Proceedings of the XXIII International Conference on Human Computer Interaction* (Association for Computing Machinery, Lleida, Spain, 2024). doi:10.1145/3612783.3612810.
140. Lam, D., Manzie, C. & Good, M. *Model predictive contouring control in Decision and Control (CDC), 2010 49th IEEE Conference on* (IEEE, 2010), 6137. doi:10.1109/CDC.2010.5717042.
141. Lam, D., Manzie, C. & Good, M. C. Model predictive contouring control for biaxial systems. *IEEE Transactions on Control Systems Technology* **21**, 552. doi:10.1109/TCST.2012.2186299 (2013).
142. Langerak, T., Christen, S., Albaba, M., Gebhardt, C., Holz, C. & Hilliges, O. *MARLUI: Multi-Agent Reinforcement Learning for Adaptive Point-and-Click UIs* in. **8** (Association for Computing Machinery, New York, NY, USA, 2024). doi:10.1145/3661147.
143. Langerak, T., Zarate, J., Vechev, V., Lindlbauer, D., Panozzo, D. & Hilliges, O. *Optimal Control for Electromagnetic Haptic Guidance Systems* in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (eds Iqbal, S. T., MacLean, K. E., Chevalier, F. & Mueller, S.) (ACM, 2020).
144. Langerak, T., Zarate, J., Vechev, V., Panozzo, D. & Hilliges, O. *A Demonstration on Dynamic Drawing Guidance via Electromagnetic Haptic Feedback* in *The Adjunct Publication of the 32nd Annual ACM Symposium on User Interface Software and Technology* (eds Guimbretièrre, F., Bernstein, M. & Reinecke, K.) (Association for Computing Machinery, New Orleans, LA, USA, 2019), 110. doi:10.1145/3332167.3356889.
145. Langerak, T., Zárate, J. J., Lindlbauer, D., Holz, C. & Hilliges, O. *Omni: Volumetric Sensing and Actuation of Passive Magnetic Tools for Dynamic Haptic Feedback* in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Association for Computing Machinery, Virtual Event, USA, 2020), 594. doi:10.1145/3379337.3415589.
146. Lashkari, Y., Metral, M. & Maes, P. Collaborative interface agents. *Readings in agents*, 111 (1997).
147. Lee, J., Wang, J., Brown, E., Chu, L., Rodriguez, S. S. & Froehlich, J. E. *GazePointAR: A Context-Aware Multimodal Voice Assistant for Pronoun Disambiguation in Wearable Augmented Reality* in *Proceedings of the CHI Conference on Human Factors in Computing Systems* (eds Mueller, F. ',

- Kyburz, P., Williamson, J. R., Sas, C., Wilson, M. L., Dugas, P. O. T. & Shklovski, I.) (ACM, 2024), 1. doi:10.48550/arxiv.2404.08213.
148. Lee, J., Sinclair, M., Gonzalez-Franco, M., Ofek, E. & Holz, C. *TORC: A Virtual Reality Controller for In-Hand High-Dexterity Finger Interaction in Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (eds Brewster, S. A., Fitzpatrick, G., Cox, A. L. & Kostakos, V.) (Association for Computing Machinery, Glasgow, Scotland Uk, 2019). doi:10.1145/3290605.3300301.
 149. Leibo, J. Z., Zambaldi, V., Lanctot, M., Marecki, J. & Graepel, T. Multi-agent reinforcement learning in sequential social dilemmas. *arXiv preprint arXiv:1702.03037*. doi:10.48550/arxiv.1702.03037 (2017).
 150. Leino, K., Todi, K., Oulasvirta, A. & Kurimo, M. *Computer-Supported Form Design Using Keystroke-Level Modeling with Reinforcement Learning in Proceedings of the 24th International Conference on Intelligent User Interfaces: Companion* (Association for Computing Machinery, Marina del Ray, California, 2019), 85. doi:10.1145/3308557.3308704.
 151. Li, J., Cho, M. K., Zhao, X., Ju, W. & Malle, B. F. *From trolley to autonomous vehicle: Perceptions of responsibility and moral norms in traffic accidents with self-driving cars in Proceedings of the Society of Automotive Engineers World Congress* (SAE International, 2016). doi:10.4271/2016-01-0164.
 152. Li, W., Liu, B., Kosasih, P. B. & Zhang, X. A 2-DOF MR actuator joystick for virtual reality applications. *Sensors and Actuators A: Physical* **137**, 308. doi:10.1016/j.sna.2007.03.015 (2007).
 153. Li, Y., Song, J. & Ermon, S. Inferring the latent structure of human decision-making from raw visual inputs. *arXiv preprint arXiv:1703.08840* **2** (2017).
 154. Liang, E., Liaw, R., Moritz, P., Nishihara, R., Fox, R., Goldberg, K., Gonzalez, J. E., Jordan, M. I. & Stoica, I. *RLlib: Abstractions for Distributed Reinforcement Learning* 2018.
 155. Liang, R.-H., Cheng, K.-Y., Su, C.-H., Weng, C.-T., Chen, B.-Y. & Yang, D.-N. *GaussSense: Attachable Stylus Sensing Using Magnetic Sensor Grid in Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (eds Miller, R., Benko, H. & Latulipe, C.) (Association for Computing Machinery, Cambridge, Massachusetts, USA, 2012), 319. doi:10.1145/2380116.2380157.

156. Liebman, E., Saar-Tsechansky, M. & Stone, P. *DJ-MC: A Reinforcement-Learning Agent for Music Playlist Recommendation in Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems* (2015), 591.
157. Lin, L.-F., Teng, S.-Y., Liang, R.-H. & Chen, B.-Y. *Stylus Assistant: Designing Dynamic Constraints for Facilitating Stylus Inputs on Portable Displays in SIGGRAPH ASIA 2016 Emerging Technologies* (Association for Computing Machinery, Macau, 2016). doi:10.1145/2988240.2988255.
158. Lindlbauer, D., Feit, A. M. & Hilliges, O. *Context-aware online adaptation of mixed reality interfaces in Proceedings of the 32nd annual ACM symposium on user interface software and technology* (eds Guimbretière, F., Bernstein, M. & Reinecke, K.) (ACM, 2019), 147. doi:10.1145/3332165.3347945.
159. Liniger, A., Domahidi, A. & Morari, M. Optimization-based autonomous racing of 1:43 scale RC cars. *Optimal Control Applications and Methods* **36**, 628. doi:10.1002/oca.2123 (2014).
160. Liu, C., Lee, S., Varnhagen, S. & Tseng, H. E. *Path planning for autonomous vehicles using model predictive control in 2017 IEEE Intelligent Vehicles Symposium (IV)* (IEEE, 2017), 174. doi:10.1109/ivs.2017.7995716.
161. Liu, F., Tang, R., Li, X., Zhang, W., Ye, Y., Chen, H., Guo, H. & Zhang, Y. Deep reinforcement learning based recommendation with explicit user-item interactions modeling. *arXiv preprint arXiv:1810.12027*. doi:10.48550/arxiv.1810.12027 (2018).
162. Liu, V., Vermeulen, J., Fitzmaurice, G. & Matejka, J. *3DALL-E: Integrating text-to-image AI in 3D design workflows in Proceedings of the 2023 ACM designing interactive systems conference* (eds Byrne, D., Martelaro, N., Boucher, A., Chatting, D. J., Alaoui, S. F., Fox, S. E., Nicenboim, I. & MacArthur, C.) (ACM, 2023), 1955. doi:10.1145/3563657.3596098.
163. Lomas, J. D., Forlizzi, J., Poonwala, N., Patel, N., Shodhan, S., Patel, K., Koedinger, K. & Brunskill, E. *Interface design optimization as a multi-armed bandit problem in Proceedings of the 2016 CHI conference on human factors in computing systems* (eds Kaye, J., Druin, A., Lampe, C., Morris, D. & Hourcade, J. P.) (ACM, 2016), 4142. doi:10.1145/2858036.2858425.
164. Long, Q., Zhou, Z., Gupta, A., Fang, F., Wu, Y. & Wang, X. Evolutionary Population Curriculum for Scaling Multi-Agent Reinforcement Learning. *CoRR abs/2003.10423* (2020).

165. Lopes, P., Yüksel, D., Guimbretière, F. & Baudisch, P. *Muscle-Plotter: An Interactive System Based on Electrical Muscle Stimulation That Produces Spatial Output* in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology* (Association for Computing Machinery, Tokyo, Japan, 2016), 207. doi:10.1145/2984511.2984530.
166. Losey, D. P., McDonald, C. G., Battaglia, E. & O'Malley, M. K. A Review of Intent Detection, Arbitration, and Communication Aspects of Shared Control for Physical Human–Robot Interaction. *Applied Mechanics Reviews* **70**, 010804. doi:10.1115/1.4039145 (2018).
167. Lowe, R., Wu, Y. I., Tamar, A., Harb, J., Pieter Abbeel, O. & Mordatch, I. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems* **30** (eds Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N. & Garnett, R.) 6379 (2017).
168. Lyons, J. B., Sycara, K., Lewis, M. & Capiola, A. Human-Autonomy Teaming: Definitions, Debates, and Directions. *Frontiers in Psychology* **12**, 589585. doi:10.3389/fpsyg.2021.589585 (2021).
169. Ma, L., Xu, N., Zhao, X., Zong, G. & Huo, X. Small-Gain Technique-Based Adaptive Neural Output-Feedback Fault-Tolerant Control of Switched Nonlinear Systems With Unmodeled Dynamics. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* **51**, 7051. doi:10.1109/TSMC.2020.2964822 (2021).
170. Ma, Z., Wang, R., Li, F.-F., Bernstein, M. & Krishna, R. Elign: Expectation alignment as a multi-agent intrinsic reward. *Advances in Neural Information Processing Systems* **35** (eds Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K. & Oh, A.) 8304. doi:10.48550/arxiv.2210.04365 (2022).
171. Mackay, W. Responding to cognitive overload: Co-adaptation between users and technology. *Intellectica* **30**, 177. doi:10.3406/intel.2000.1597 (2000).
172. Madni, A. M. & Madni, C. C. Architectural framework for exploring adaptive human-machine teaming options in simulated dynamic environments. *Systems* **6**, 44. doi:10.3390/systems6040044 (2018).
173. Maes, P. in *Readings in human–computer interaction* 811 (Elsevier, 1995). doi:10.1016/b978-0-08-051574-8.50084-4.

174. Maqueda, A. I., Loquercio, A., Gallego, G., García, N. & Scaramuzza, D. *Event-based vision meets deep learning on steering prediction for self-driving cars* in *Proceedings of the IEEE conference on computer vision and pattern recognition* (IEEE, 2018), 5419. doi:10.1109/cvpr.2018.00568.
175. Marcano, M., Díaz, S., Pérez, J. & Irigoyen, E. A review of shared control for automated vehicles: Theory and applications. *IEEE Transactions on Human-Machine Systems* **50**, 475. doi:10.1109/thms.2020.3017748 (2020).
176. Martín, J. A. Á., Gollee, H., Müller, J. & Murray-Smith, R. Intermittent control as a model of mouse movements. *ACM Transactions on Computer-Human Interaction (TOCHI)* **28**, 1. doi:10.1145/3461836 (2021).
177. Massie, T. H. & Salisbury, J. K. *The PHANToM Haptic Interface: A Device for Probing Virtual Objects* in *Proceedings of the ASME Winter Annual Meeting '94* (1994), 295.
178. McCreath, E., Kay, J. & Crawford, E. IEMS-an approach that combines handcrafted rules with learnt instance based rules. *Aust. J. Intell. Inf. Process. Syst.* **9**, 40 (2006).
179. McIntosh, J., Strohmeier, P., Knibbe, J., Boring, S. & Hornbæk, K. *Magnetips: Combining Fingertip Tracking and Haptic Feedback for Around-Device Interaction* in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (eds Brewster, S. A., Fitzpatrick, G., Cox, A. L. & Kostakos, V.) (ACM, Glasgow, Scotland Uk, 2019), 408:1. doi:10.1145/3290605.3300638.
180. McRuer, D. T. & Jex, H. R. A review of quasi-linear pilot models. *IEEE transactions on human factors in electronics* **HFE-8**, 231. doi:10.1109/thfe.1967.234304 (1967).
181. Mehrotra, A., Musolesi, M., Hendley, R. & Pejovic, V. *Designing content-driven intelligent notification mechanisms for mobile applications* in *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (ACM, 2015), 813. doi:10.1145/2750858.2807544.
182. Miller, C. A. & Parasuraman, R. Beyond Levels of Automation: An Architecture for More Flexible Human-Automation Collaboration. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* **47**, 182. doi:10.1177/154193120304700138 (2003).
183. Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D. & Riedmiller, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*. doi:10.48550/arxiv.1312.5602 (2013).

184. Moerland, T. M., Broekens, J., Plaat, A., Jonker, C. M., *et al.* Model-based reinforcement learning: A survey. *Foundations and Trends® in Machine Learning* **16**, 1. doi:10.1561/22000000086 (2023).
185. morph. *Sensel Morph Official Site* <https://sensel.com/pages/the-sensel-morph>. 2017.
186. Mou, Y. & Xu, K. The media inequality: Comparing the initial human-human and human-AI social interactions. *Computers in Human Behavior* **72**, 432. doi:10.1016/j.chb.2017.02.067 (2017).
187. Mueller, M. & D'Andrea, R. *A model predictive controller for quadcopter state interception in Proceedings of the European Control Conference (ECC), 2013* (IEEE, 2013), 1383. doi:10.23919/ECC.2013.6669415.
188. Müller, J., Oulasvirta, A. & Murray-Smith, R. Control theoretic models of pointing. *ACM Transactions on Computer-Human Interaction (TOCHI)* **24**, 1. doi:10.1145/3121431 (2017).
189. Mullins, J., Mawson, C. & Nahavandi, S. *Haptic handwriting aid for training and rehabilitation in Systems, Man and Cybernetics, 2005 IEEE International Conference on* **3** (IEEE, 2005), 2690. doi:10.1109/ICSMC.2005.1571556.
190. Murray-Smith, R. Control Theory, Dynamics, and Continuous Interaction. *Computational Interaction*. doi:10.1093/oso/9780198799603.003.0002 (2018).
191. Murray-Smith, R., Oulasvirta, A., Howes, A., Müller, J., Ikkala, A., Bachinski, M., Fleig, A., Fischer, F. & Klar, M. What simulation can do for HCI research. *Interactions* **29**, 48. doi:10.1145/3564038 (2022).
192. Murray-Smith, R., Williamson, J. H., Ramsay, A., Tonolini, F., Rogers, S. & Loriette, A. Forward and inverse models in HCI: physical simulation and deep learning for inferring 3D finger pose. *arXiv preprint arXiv:2109.03366 abs/2109.03366* (2021).
193. Nägeli, T., Meier, L., Domahidi, A., Alonso-Mora, J. & Hilliges, O. *Real-time Planning for Automated Multi-View Drone Cinematography* in. **36** (Association for Computing Machinery (ACM), Los Angeles, USA, 2017), 1.
194. Nakagaki, K. & Kakehi, Y. *Comp*Pass: A Compass-Based Drawing Interface in CHI '14 Extended Abstracts on Human Factors in Computing Systems* (eds Jones, M., Palanque, P. A., Schmidt, A. & Grossman, T.) (Association for Computing Machinery, Toronto, Ontario, Canada, 2014), 447. doi:10.1145/2559206.2574766.
195. NeuroDigital Technologies. *Gloveone Glove* <https://www.neurodigital.es/gloveone/>. Last accessed: 29.03.2017.

196. O'Neill, T., McNeese, N., Barron, A. & Schelble, B. Human–Autonomy Teaming: A Review and Analysis of the Empirical Literature. *Human Factors* **64**. PMID: 33092417, 904. doi:10.1177/0018720820960865 (2022).
197. Ogata, M. *Magneto-Haptics: Embedding Magnetic Force Feedback for Physical Interactions in The 31st Annual ACM Symposium on User Interface Software and Technology* (eds Baudisch, P., Schmidt, A. & Wilson, A.) (ACM, 2018), 737. doi:10.1145/3242587.3242615.
198. OpenAI. *ChatGPT: The GPT-4 Language Model* <https://www.openai.com/chatgpt>. 2023.
199. Orr, J. & Dutta, A. Multi-Agent Deep Reinforcement Learning for Multi-Robot Applications: A Survey. *Sensors* **23**. doi:10.3390/s23073625 (2023).
200. Ota, J. Multi-agent robot systems as distributed autonomous systems. *Advanced Engineering Informatics* **20**, 59. doi:<https://doi.org/10.1016/j.aei.2005.06.002> (2006).
201. Oulasvirta, A., Dayama, N. R., Shiripour, M., John, M. & Karrenbauer, A. Combinatorial Optimization of Graphical User Interface Designs. *Proceedings of the IEEE* **108**, 434. doi:10.1109/JPROC.2020.2969687 (2020).
202. Oulasvirta, A., Feit, A., Lähenteenlahti, P. & Karrenbauer, A. Computational Support for Functionality Selection in Interaction Design. *ACM Transactions on Computer-Human Interaction* **24**, 1. doi:10.1145/3131608 (2017).
203. Oulasvirta, A., Jokinen, J. P. & Howes, A. *Computational Rationality as a Theory of Interaction in CHI Conference on Human Factors in Computing Systems* (ACM, 2022), 1. doi:10.1145/3491102.3517739.
204. Oulasvirta, A., Kristensson, P. O., Bi, X. & Howes, A. *Computational interaction* doi:10.1145/3170427.3170654 (Oxford University Press, 2018).
205. Pangaro, G., Maynes-Aminzade, D. & Ishii, H. *The Actuated Workbench: Computer-Controlled Actuation in Tabletop Tangible Interfaces in Proceedings of the 15th Annual ACM Symposium on User Interface Software and Technology* (ed Beaudouin-Lafon, M.) (Association for Computing Machinery, Paris, France, 2002), 181. doi:10.1145/571985.572011.
206. Park, S., Gebhardt, C., Rädle, R., Feit, A. M., Vrzakova, H., Dayama, N. R., Yeo, H.-S., Klokmoose, C. N., Quigley, A., Oulasvirta, A., *et al.* *Adam: Adapting multi-user interfaces for collaborative environments in real-time in Proceedings of the 2018 CHI conference on human factors in computing systems* (2018), 1. doi:10.1145/3173574.3173758.

207. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. & Duchesnay, E. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **12**, 2825 (2011).
208. Pejovic, V. & Musolesi, M. *InterruptMe: designing intelligent prompting mechanisms for pervasive applications* in *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (eds Brush, A. J., Friday, A., Kientz, J. A., Scott, J. & Song, J.) (ACM, 2014), 897. doi:10.1145/2632048.2632062.
209. Peng, H., Zoran, A. & Guimbreti re, F. V. *D-Coil: A Hands-on Approach to Digital 3D Models Design* in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (eds Begole, B., Kim, J., Inkpen, K. & Woo, W.) (Association for Computing Machinery, Seoul, Republic of Korea, 2015), 1807. doi:10.1145/2702123.2702381.
210. PerplexityAI. *Perplexity AI: An AI-powered Search Engine* <https://www.perplexity.ai>.
211. Petruska, A. J. & Abbott, J. J. Omnimagnet: An omnidirectional electromagnet for controlled dipole-field generation. *IEEE Transactions on Magnetics* **50**, 1. doi:10.1109/tmag.2014.2303784 (2014).
212. Pisharoty, D., Yang, R., Newman, M. W. & Whitehouse, K. *Thermocoach: Reducing home energy consumption with personalized thermostat recommendations* in *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments* (eds Culler, D., Agarwal, Y. & Mangharam, R.) (ACM, 2015), 201. doi:10.1145/2821650.2821671.
213. Polydoros, A. S. & Nalpantidis, L. Survey of model-based reinforcement learning: Applications on robotics. *Journal of Intelligent & Robotic Systems* **86**, 153. doi:10.1007/s10846-017-0468-y (2017).
214. Pothi, N. Design and Control of One Degree of Freedom Haptic Interface System with PID Controller. *Advanced Materials Research* **931**, 1323. doi:10.4028/www.scientific.net/amr.931-932.1323 (2014).
215. Poulton, E. C. *Tracking Skill and Manual Control* (Academic Press, New York, 1974).
216. Prada, R. & Paiva, A. *Human-agent interaction: Challenges for bringing humans and agents together* in *Proc. of the 3rd Int. Workshop on Human-Agent Interaction Design and Models (HAIDM 2014) at the 13th Int. Conf. on Agent and Multi-Agent Systems (AAMAS 2014)* (2014), 1.

217. Prattichizzo, D., Chinello, F., Pacchierotti, C. & Malvezzi, M. Towards Wearability in Fingertip Haptics: a 3-DoF Wearable Device for Cutaneous Force Feedback. *EEE Trans. Haptics* **6**, 506. doi:10.1109/TOH.2013.53 (2013).
218. Qin, S. & Badgwell, T. A. A survey of industrial model predictive control technology. *Control Engineering Practice* **11**, 733. doi:https://doi.org/10.1016/S0967-0661(02)00186-7 (2003).
219. Ramos-Velasco, L. E., Domínguez-Ramírez, O. A. & Parra-Vega, V. Wavenet fuzzy PID controller for nonlinear MIMO systems: Experimental validation on a high-end haptic robotic interface. *Applied Soft Computing* **40**, 199. doi:10.1016/j.asoc.2015.11.014 (2016).
220. Reilly, D., Welsman-Dinelle, M., Bate, C. & Inkpen, K. *Just point and click? using handhelds to interact with paper maps in Proceedings of the 7th International Conference on Human Computer Interaction with Mobile Devices & Services* (Association for Computing Machinery, Salzburg, Austria, 2005), 239. doi:10.1145/1085777.1085820.
221. Rich, C. & Sidner, C. L. in *Computational Models of Mixed-Initiative Interaction* 149 (Springer, 1998). doi:10.1007/978-94-017-1118-0_4.
222. Rich, C., Sidner, C., Lesh, N., Garland, A., Booth, S. & Chimani, M. *DiamondHelp: A collaborative interface framework for networked home appliances in 25th IEEE International Conference on Distributed Computing Systems Workshops* (IEEE, 2005), 514. doi:10.1109/icdcs.2005.58.
223. Rizzoglio, F., Casadio, M., De Santis, D. & Mussa-Ivaldi, F. A. Building an adaptive interface via unsupervised tracking of latent manifolds. *Neural Networks* **137**, 174. doi:https://doi.org/10.1016/j.neunet.2021.01.009 (2021).
224. Rockafellar, R. T. & Wets, R. J.-B. *Variational analysis* doi:10.1007/978-3-642-02431-3 (Springer Science & Business Media, 2009).
225. Rodriguez-Soto, M., Serramia, M., Lopez-Sanchez, M. & Rodriguez-Aguilar, J. A. Instilling moral value alignment by means of multi-objective reinforcement learning. *Ethics and Information Technology* **24**, 9. doi:10.1007/s10676-022-09635-0 (2022).
226. Russel, M. *A neural network driven sensor array for locating a permanent magnet* MA thesis (Graduate Studies, 2017). doi:10.11575/PRISM/24856.
227. Salvucci, D. D. An integrated model of eye movements and visual encoding. *Cognitive Systems Research* **1**, 201. doi:10.1016/s1389-0417(00)00015-2 (2001).

228. Schlageter, V., Besse, P.-A., Popovic, R. & Kucera, P. Tracking system with five degrees of freedom using a 2D-array of Hall sensors and a permanent magnet. *Sensors and Actuators A: Physical* **92**, 37. doi:10.1016/S0924-4247(01)00537-4 (2001).
229. Schuetz, I., Murdison, T. S., MacKenzie, K. J. & Zannoli, M. *An Explanation of Fitts' Law-like Performance in Gaze-Based Selection Tasks Using a Psychophysics Approach in Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Association for Computing Machinery, Glasgow, Scotland Uk, 2019), 1. doi:10.1145/3290605.3300765.
230. Schulman, J., Levine, S., Abbeel, P., Jordan, M. & Moritz, P. *Trust region policy optimization in International conference on machine learning abs/1502.05477* (2015), 1889.
231. Schulman, J., Wolski, F., Dhariwal, P., Radford, A. & Klimov, O. *Proximal Policy Optimization Algorithms* 2017. doi:10.48550/arxiv.1707.06347.
232. Senkal, D. & Gurocak, H. Spherical brake with MR fluid as multi degree of freedom actuator for haptics. *Journal of Intelligent Material Systems and Structures* **20**, 2149. doi:10.1177/1045389X09348925 (2009).
233. Senkal, D. & Gurocak, H. Haptic joystick with hybrid actuator using air muscles and spherical MR-brake. *Mechatronics* **21**, 951. doi:10.1016/j.mechatronics.2011.03.001 (2011).
234. Seo, Y.-W. & Zhang, B.-T. *A reinforcement learning agent for personalized information filtering in Proceedings of the 5th international conference on Intelligent user interfaces* (ACM, 2000), 248. doi:10.1145/325737.325859.
235. Sethi, S. *Optimal control theory: Applications to management science and economics* doi:10.1007/978-3-030-91745-6 (Springer, 2002).
236. Shahriari, B., Swersky, K., Wang, Z., Adams, R. P. & de Freitas, N. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proceedings of the IEEE* **104**, 148. doi:10.1109/JPROC.2015.2494218 (2016).
237. Shapley, L. S. Stochastic games. *Proceedings of the national academy of sciences* **39**, 1095 (1953).
238. Shen, J., Fitzhenry, E. & Dietterich, T. G. *Discovering frequent work procedures from resource connections in Proceedings of the 14th international conference on Intelligent user interfaces* (eds Conati, C., Bauer, M., Oliver, N. & Weld, D. S.) (ACM, 2009), 277. doi:10.1145/1502650.1502690.

239. Shen, J., Irvine, J., Bao, X., Goodman, M., Kolibaba, S., Tran, A., Carl, F., Kirschner, B., Stumpf, S. & Dietterich, T. G. *Detecting and correcting user activity switches: algorithms and interfaces* in *Proceedings of the 14th International Conference on Intelligent User Interfaces* (Association for Computing Machinery, Sanibel Island, Florida, USA, 2009), 117. doi:10.1145/1502650.1502670.
240. Sheridan, T. B. Adaptive automation, level of automation, allocation authority, supervisory control, and adaptive control: distinctions and modes of adaptation. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* **41**, 662. doi:10.1109/tsmca.2010.2093888 (2011).
241. Sheridan, T. B. & Ferrell, W. R. *Man-machine systems; Information, control, and decision models of human performance*. doi:10.1115/1.3426860 (the MIT press, 1974).
242. Sheridan, T. B. & Verplank, W. L. *Human and computer control of undersea teleoperators* in *The Human Operator in Complex Systems* (Defense Technical Information Center, 1978).
243. Shilkrot, R., Maes, P. & Zoran, A. *Physical Rendering with a Digital Airbrush* in *ACM SIGGRAPH 2014 Studio* (Association for Computing Machinery, Vancouver, Canada, 2014), 40:1. doi:10.1145/2619195.2656328.
244. Shively, R. J., Lachter, J., Brandt, S. L., Matessa, M., Battiste, V. & Johnson, W. W. *A human-autonomy teaming approach for a flight-following task* in *International Conference on Applied Human Factors and Ergonomics* (Springer Nature, 2017), 12. doi:10.1007/978-3-319-60642-2_2.
245. Shneiderman, B. & Maes, P. Direct manipulation vs. interface agents. *interactions* **4**, 42. doi:10.1145/267505.267514 (1997).
246. Simić, S. N. A short introduction to geometric control theory. *arXiv preprint arXiv:2404.00059*. doi:10.48550/arXiv.2404.00059 (2024).
247. Simo-Serra, E., Iizuka, S. & Ishikawa, H. Mastering Sketching: Adversarial Augmentation for Structured Prediction. *ACM Trans. Graph.* **37**, 1. doi:10.1145/3132703 (2018).
248. Simo-Serra, E., Iizuka, S., Sasaki, K. & Ishikawa, H. Learning to Simplify: Fully Convolutional Networks for Rough Sketch Cleanup. *ACM Trans. Graph.* **35**, 1. doi:10.1145/2897824.2925972 (2016).

249. Sinclair, M., Ofek, E., Gonzalez-Franco, M. & Holz, C. *CapstanCrunch: A Haptic VR Controller with User-Supplied Force Feedback* in *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology* (eds Guimbretière, F., Bernstein, M. & Reinecke, K.) (Association for Computing Machinery, New Orleans, LA, USA, 2019), 815. doi:10.1145/3332165.3347891.
250. Smith, D. A. & Lieberman, H. *The why UI: using goal networks to improve user interfaces* in *Proceedings of the 15th international conference on Intelligent user interfaces* (2010), 377.
251. Soh, H., Sanner, S., White, M. & Jamieson, G. *Deep sequential recommendation for personalized adaptive user interfaces* in *Proceedings of the 22nd international conference on intelligent user interfaces* (eds Papadopoulos, G. A., Kuflik, T., Chen, F., Duarte, C. & Fu, W.-T.) (ACM, 2017), 589. doi:10.1145/3025171.3025207.
252. Somefun, O. A., Akingbade, K. & Dahunsi, F. The dilemma of PID tuning. *Annual Reviews in Control* **52**, 65. doi:10.1016/j.arcontrol.2021.05.002 (2021).
253. Stamper, R. E. *A Three Degree Of Freedom Parallel Manipulator With Only Translational Degrees Of Freedom* PhD thesis (University of Maryland at College Park, 1997).
254. Stephanidis, C., Karagiannidis, C. & Koumpis, A. *Decision making in intelligent user interfaces* in *Proceedings of the 2nd international conference on Intelligent user interfaces* (ACM Press, 1997), 195. doi:10.1145/238218.238323.
255. Strasnick, E., Yang, J., Tanner, K., Olwal, A. & Follmer, S. *ShiftIO: Reconfigurable Tactile Elements for Dynamic Affordances and Mobile Interaction* in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (eds Mark, G., Fussell, S. R., Lampe, C., schraefel m. c., m. c., Hourcade, J. P., Appert, C. & Wigdor, D.) (Association for Computing Machinery, Denver, Colorado, USA, 2017), 5075. doi:10.1145/3025453.3025988.
256. Su, P.-H., Budzianowski, P., Ultes, S., Gasic, M. & Young, S. Sample-efficient actor-critic reinforcement learning with supervised data for dialogue management. *arXiv preprint arXiv:1707.00130* **abs/1707.00130** (eds Jokinen, K., Stede, M., DeVault, D. & Louis, A.) doi:10.48550/arxiv.1707.00130 (2017).
257. Sutton, R. S. & Barto, A. G. *Reinforcement Learning: An Introduction* (A Bradford Book, Cambridge, MA, USA, 2018).

258. Suzuki, R., Kato, J., Gross, M. D. & Yeh, T. *Reactile: Programming Swarm User Interfaces through Direct Physical Manipulation* in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (eds Mandryk, R. L., Hancock, M., Perry, M. & Cox, A. L.) (Association for Computing Machinery, Montreal QC, Canada, 2018), 1. doi:10.1145/3173574.3173773.
259. Sycara, K. & Lewis, M. in *Team cognition: Understanding the factors that drive process and performance 203* (American Psychological Association, 2004). doi:10.1037/10690-000.
260. Tan, M. *Multi-agent reinforcement learning: Independent vs. cooperative agents* in *Proceedings of the tenth international conference on machine learning* (Elsevier BV, 1993), 330. doi:10.1016/b978-1-55860-307-3.50049-6.
261. Tarantino, S., Clemente, F., Barone, D., Controzzi, M. & Cipriani, C. The myokinetic control interface: tracking implanted magnets as a means for prosthetic control. *Scientific reports* **7**, 1. doi:10.1038/s41598-017-17464-1 (2017).
262. Taylor, C. R., Abramson, H. G. & Herr, H. M. Low-Latency Tracking of Multiple Permanent Magnets. *IEEE Sensors Journal* **19**, 11458. doi:10.1109/JSEN.2019.2936766 (2019).
263. Teranishi, A., Korres, G., Park, W. & Eid, M. Combining full and partial haptic guidance improves handwriting skills development. *IEEE transactions on haptics* **11**, 509. doi:10.1109/TOH.2018.2851511 (2018).
264. Thomaszewski, B., Gumann, A., Pabst, S. & Straßer, W. *Magnets in motion* in *ACM Transactions on Graphics (TOG)* **27** (Association for Computing Machinery (ACM), 2008), 162. doi:10.1145/1409060.1409115.
265. Tian, Z., Zou, S., Davies, I., Warr, T., Wu, L., Ammar, H. B. & Wang, J. *Learning to communicate implicitly by actions* in *Proceedings of the AAAI Conference on Artificial Intelligence* **34** (Association for the Advancement of Artificial Intelligence (AAAI), 2020), 7261. doi:10.1609/aaai.v34i05.6217.
266. Todi, K., Bailly, G., Leiva, L. & Oulasvirta, A. *Adapting User Interfaces with Model-based Reinforcement Learning* in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)* (eds Kitamura, Y., Quigley, A., Isbister, K., Igarashi, T., Bjørn, P. & Drucker, S. M.) (ACM, 2021). doi:10.1145/3411764.3445497.

267. Van Der Linde, R. Q., Lammertse, P., E., F. & B., R. *The HapticMaster, a New Highperformance Haptic Interface in Proceedings of Eurohaptics 2002* (2002).
268. Wang, W., Yang, T., Liu, Y., Hao, J., Hao, X., Hu, Y., Chen, Y., Fan, C. & Gao, Y. *From few to more: Large-scale dynamic multiagent curriculum learning in Proceedings of the AAAI Conference on Artificial Intelligence* **34** (Association for the Advancement of Artificial Intelligence (AAAI), 2020), 7293. doi:10.1609/aaai.v34i05.6221.
269. Wei Xu Marvin J. Dainoff, L. G. & Gao, Z. Transitioning to Human Interaction with AI Systems: New Challenges and Opportunities for HCI Professionals to Enable Human-Centered AI. *International Journal of Human-Computer Interaction* **39**, 494. doi:10.1080/10447318.2022.2041900 (2023).
270. Weiss, M., Schwarz, F., Jakubowski, S. & Borchers, J. *Madgets: Actuating Widgets on Interactive Tabletops in Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology* (eds Perlin, K., Czerwinski, M. & Miller, R.) (Association for Computing Machinery, New York, New York, USA, 2010), 293. doi:10.1145/18666029.1866075.
271. Weiss, M., Wacharamanotham, C., Voelker, S. & Borchers, J. *FingerFlux: near-surface haptic feedback on tabletops in Proceedings of the 24th annual ACM symposium on User interface software and technology* (2011), 615. doi:10.1145/2047196.2047277.
272. Wellman, P. & Howe, R. D. *Towards Realistic Vibrotactile Display in Virtual Environments in Proceedings of the ASME Dynamic Systems and Control Division* (1995), 713.
273. Wermelinger, M. *Using GitHub Copilot to Solve Simple Programming Problems in Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. I* (Association for Computing Machinery, Toronto ON, Canada, 2023), 172. doi:10.1145/3545945.3569830.
274. Whitmire, E., Benko, H., Holz, C., Ofek, E. & Sinclair, M. *Haptic revolver: Touch, shear, texture, and shape rendering on a reconfigurable virtual reality controller in Proceedings of the 2018 CHI conference on human factors in computing systems* (eds Mandryk, R. L., Hancock, M., Perry, M. & Cox, A. L.) (ACM, 2018), 1. doi:10.1145/3173574.3173660.
275. Wickens, C. D. & Hollands, J. G. *Engineering psychology and human performance* 3rd ed. (Prentice-Hall, New York, NY, 1999).

276. Wiener, N. *Cybernetics: Control and communication in the animal and the machine* (MIT Press, Cambridge, MA, 1948).
277. Xing, J., Wei, L.-Y., Shiratori, T. & Yatani, K. Autocomplete Hand-Drawn Animations. *ACM Trans. Graph.* **34**, 1. doi:10.1145/2816795.2818079 (2015).
278. Yamaoka, J. & Kakehi, Y. *dePENd: augmented handwriting system using ferromagnetism of a ballpoint pen* in *Proceedings of the 26th annual ACM symposium on User interface software and technology* (2013), 203.
279. Yang, R. & Newman, M. Learning from a learning thermostat: Lessons for intelligent systems for the home. *UbiComp 2013 - Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 93. doi:10.1145/2493432.2493489 (2013).
280. Yang, Z., Huang, L., Chen, Y., Wei, Z., Ahn, S., Zelinsky, G., Samaras, D. & Hoai, M. *Predicting goal-directed human attention using inverse reinforcement learning* in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (IEEE, 2020), 193. doi:10.1109/cvpr42600.2020.00027.
281. Yorke-Smith, N., Saadati, S., Myers, K. L. & Morley, D. N. The design of a proactive personal agent for task management. *International Journal on Artificial Intelligence Tools* **21**, 1250004 (2012).
282. Yoshida, S., Noma, H. & Hosaka, K. *Proactive desk II: Development of a new multi-object haptic display using a linear induction motor* in *Virtual Reality Conference, 2006* (IEEE Computer Society, 2006), 269. doi:10.1109/VR.2006.110.
283. Yu, C., Velu, A., Vinitsky, E., Gao, J., Wang, Y., Bayen, A. & Wu, Y. The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in Neural Information Processing Systems* **35** (eds Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K. & Oh, A.) 24611 (2022).
284. Yu, C., Zhang, M., Ren, F. & Tan, G. Emotional Multiagent Reinforcement Learning in Spatial Social Dilemmas. *IEEE Transactions on Neural Networks and Learning Systems* **26**, 3083. doi:10.1109/TNNLS.2015.2403394 (2015).
285. Yung, K. W., Landecker, P. B. & Villani, D. D. An analytic solution for the force between two magnetic dipoles. *Physical Separation in Science and Engineering* **9**, 39. doi:10.1155/1998/79537 (1998).

286. Zarate, J. J., Langerak, T., Thomaszewski, B. & Hilliges, O. Contact-free Nonplanar Haptics with a Spherical Electromagnet. *IEEE Haptics Symposium (HAPTICS)*, 698. doi:10.1109/HAPTICS45997.2020.ras.HAP20.33.99093c10 (2020).
287. Zhang, K., Yang, Z. & Başar, T. in *Handbook of Reinforcement Learning and Control* (eds Vamvoudakis, K. G., Wan, Y., Lewis, F. L. & Cansever, D.) 321 (Springer International Publishing, Cham, 2021). doi:10.1007/978-3-030-60990-0_12.
288. Zhang, Q., Dong, H. & El Saddik, A. Magnetic field control for haptic display: system design and simulation. *IEEE Access* **4**, 299. doi:10.1109/access.2016.2514978 (2016).
289. Zhang, R., Boyles, A. J. & Abbott, J. J. *Six principal modes of vibrotactile display via stylus* in *2018 IEEE Haptics Symposium (HAPTICS)* (eds Kuchenbecker, K. J., Gerling, G. J. & Visell, Y.) (IEEE, 2018), 313. doi:10.1109/haptics.2018.8357194.
290. Zheng, C., Kim, J., Leithinger, D., Gross, M. D. & Do, E. Y.-L. *Mechamagnets: Designing and Fabricating Haptic and Functional Physical Inputs with Embedded Magnets* in *Proceedings of the Thirteenth International Conference on Tangible, Embedded, and Embodied Interaction* (Association for Computing Machinery, Tempe, Arizona, USA, 2019), 325. doi:10.1145/3294109.3295622.
291. Zoller, E. I., Cattin, P. C., Zam, A. & Rauter, G. *Assessment of the Functional Rotational Workspace of Different Grasp Type Handles for the lambda. 6 Haptic Device* in *2019 IEEE World Haptics Conference (WHC)* (IEEE, 2019), 127. doi:10.1109/whc.2019.8816080.
292. Zoran, A. & Paradiso, J. A. *FreeD: A Freehand Digital Sculpting Tool* in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (eds Mackay, W. E., Brewster, S. A. & Bødker, S.) (Association for Computing Machinery, Paris, France, 2013), 2613. doi:10.1145/2470654.2481361.
293. Zoran, A., Shilkrot, R., Goyal, P., Maes, P. & Paradiso, J. A. The wise chisel: The rise of the smart handheld tool. *IEEE Pervasive Computing* **13**, 48. doi:10.1109/MPRV.2014.59 (2014).

GENERATIVE AI USED

The following tools have been used during the writing of this dissertation. Prompts were, or in similar spirit to, "Correct the grammar and syntax [paragraph]", or "Give

feedback on [paragraph]." Furthermore, they have been used to translate the English abstract into German.

1. OpenAI. *ChatGPT: The GPT-4 Language Model* <https://www.openai.com/chatgpt>. 2023
2. PerplexityAI. *Perplexity AI: An AI-powered Search Engine* <https://www.perplexity.ai>

Part V

APPENDIX

OPTIMAL CONTROL FOR ELECTROMAGNETIC HAPTIC GUIDANCE SYSTEMS

A.1 DIPOLE-DIPOLE MODEL

A.1.1 *In-Plane*

In this section we describe the derivation of the dipole-dipole model for the in-plane actuation force, as well as the case of considering a pen tilt θ of the pen. Please refer to the schematic Figure 8.6 for vector notations we use in this section.

The coordinate system is given by,

$$\mathbf{e}_d = \frac{\mathbf{p}_e - \mathbf{p}_p}{\|\mathbf{p}_e - \mathbf{p}_p\|} \quad (\text{A.1})$$

$$\mathbf{e}_z = [0, 0, 1]^T \quad (\text{A.2})$$

$$\mathbf{e}_t = \mathbf{e}_d \times \mathbf{e}_z \quad (\text{A.3})$$

with \mathbf{e}_d the in-paper-plane distance from the pen contact point to the electromagnet center projection, \mathbf{e}_z the vertical out-of-plane direction and \mathbf{e}_t the orthogonal vector to the former two.

The dipole-dipole expression for the force acting on \mathbf{m}_p due to \mathbf{m}_e and separated by \mathbf{r}_p is:

$$\mathbf{F}_p = \frac{3\mu_0}{4\pi r_p^5} [(\langle \mathbf{m}_p, \mathbf{r}_p \rangle) \mathbf{m}_e + (\langle \mathbf{m}_e, \mathbf{r}_p \rangle) \mathbf{m}_p + (\langle \mathbf{m}_p, \mathbf{m}_e \rangle) \mathbf{r}_p - \frac{5(\langle \mathbf{m}_p, \mathbf{r}_p \rangle)(\langle \mathbf{m}_e, \mathbf{r}_p \rangle)}{r_p^2} \mathbf{r}_p] , \quad (\text{A.4})$$

The two dipoles and the vector distance between them can be expressed in the proposed coordinate system as,

$$\mathbf{m}_e = \alpha m_m \mathbf{e}_z \quad (\text{A.5})$$

$$\begin{aligned} \mathbf{m}_p &= -(m_p \sin \theta \cos \varphi) \mathbf{e}_d \\ &\quad + (m_p \sin \theta \sin \varphi) \mathbf{e}_t \\ &\quad + (m_p \cos \theta) \mathbf{e}_z \end{aligned} \quad (\text{A.6})$$

$$\begin{aligned} \mathbf{r}_p &= -(d + h_p \sin \theta \cos \varphi) \mathbf{e}_d \\ &\quad + (h_p \sin \theta \sin \varphi) \mathbf{e}_t \\ &\quad + (h - (1 - \cos \theta)h_p) \mathbf{e}_z \end{aligned} \quad (\text{A.7})$$

and the three scalar products of equation A.4,

$$\langle \mathbf{m}_e, \mathbf{r}_p \rangle = \alpha m_m [h - (1 - \cos \theta)h_p] \quad (\text{A.8})$$

$$\begin{aligned} \langle \mathbf{m}_p, \mathbf{r}_p \rangle &= m_p [-\sin \theta \cos \varphi (d + h_p \sin \theta \cos \varphi) \\ &\quad + \sin^2 \theta \sin^2 \varphi h_p + \\ &\quad \cos \theta (h - h_p(1 - \cos \theta))] \end{aligned} \quad (\text{A.9})$$

$$\langle \mathbf{m}_e, \mathbf{m}_p \rangle = \alpha m_m m_p \cos \theta \quad (\text{A.10})$$

A.1.1.1 Position-aware dipole-dipole model

We first derive the position-aware dipole-dipole model, before continuing to the full position-aware and angle-aware model. We rewrite Eq. A.4 with an equivalent pen dipole $\tilde{\mathbf{m}}_p$, obtained by applying the small tilting angle approximation ($\cos \theta \simeq 1$ and $\sin \theta \simeq 0$) to Eq. A.6,

$$\tilde{\mathbf{m}}_p = m_p \mathbf{e}_z, \quad (\text{A.11})$$

where the scalar magnetization is given by $m_p = B_r V / \mu_0$. B_r is the residual magnetization of the permanent magnet and V its volume and \mathbf{e}_z is the z -unit vector. This approximation removes the requirement for tracking the pen tilt. More importantly it drastically simplifies the force equation since both dipoles now only have a z component and thus the actuation only depends on the distance d between pen and magnet (not on θ nor φ). This provides a simplified version of the 3D distance vector,

$$\tilde{\mathbf{r}}_p = -d \mathbf{e}_d + h \mathbf{e}_z, \quad (\text{A.12})$$

where the vertical distance, $h = h_m + h_p$, is constant. Note that the in-plane distance $d = \|\mathbf{p}_p - \mathbf{p}_e\|$ is one of the variables we seek to control, given the projections of the pen position (\mathbf{p}_e) and the electromagnet position (\mathbf{p}_p) onto the sketching plane.

Name	Value	Description
μ_0	$4\pi \cdot 10^{-7}$ [H/m]	Vacuum permeability
B_r	1.3 [T]	Pen magnet type (NIB N42)
V	0.66 [cm ³]	Pen magnet volume
m_p	0.683 [A m ²]	pen dipole (= $B_r V / \mu_0$)
m_m	1.286 [A m ²]	electromagnet dipole
h	2.71 [cm]	z-distance \mathbf{m}_e to \mathbf{m}_p
h_p	1.40 [cm]	height pen-tip to magnet
h_m	1.31 [cm]	z-distance from plane to \mathbf{m}_e .
F_0	0.488 [N]	force factor in Eq. A.14

TABLE A.1: List of electromagnet model and hardware parameters

The electromagnet dipole (\mathbf{m}_e) is mounted in a fixed upright position. Therefore it can be expressed via Eq. A.5, without incurring any approximation error. The magnetization value of the full-strength dipole m_m , which approximates the electromagnet, can be derived experimentally. For this purpose we scan the magnetic field generated by the electromagnet, setting $\alpha = 1$ and using a hall sensor and adjust the parameters of EM field equation to give a good fit, as explained below in section *Electromagnet dipole equivalent*. Table A.1 reports the values of m_m , m_p and h that were used in our experiments.

The total force acting on the pen (Eq. A.4) can now be decomposed into the in-plane and vertical force components:

$$\mathbf{F}_p = F_a \mathbf{e}_d + F_z \mathbf{e}_z. \quad (\text{A.13})$$

Here $\mathbf{F}_a = F_a \mathbf{e}_d$ represents the quantity we seek to control. By substituting the results from Eq. A.5, A.11 and A.12 into Eq. A.4 and maintaining only the in-plane contributions (\mathbf{e}_d direction), we obtain the expression for the actuation force as function of pen-magnet separation:

$$\mathbf{F}_a = \alpha F_0 \left(\frac{d \left(4 - \frac{d^2}{h^2} \right)}{h \left(1 + \frac{d^2}{h^2} \right)^{\frac{7}{2}}} \right) \mathbf{e}_d. \quad (\text{A.14})$$

where F_0 is a constant force parameter given by the expression,

$$F_0 = \frac{3 \mu_0 m_p m_m}{4 \pi h^4}. \quad (\text{A.15})$$




FIGURE A.1: In-plane magnetic force as function of position. The horizontal displacement between curves (each denoting a different pen-tilt) is the approximation error induced by the upright pen (purple) assumption (angles defined in Figure 8.6).

Fig. A.1 illustrates how the dimensionless ratio within parentheses in Eq. A.14 governs the force strength as function of distance $d = \|\mathbf{r}_d\|$. The actuation force F_a is zero if the two magnets are aligned with one another ($d = 0$), it has a maximum $F_a^{max} = 0.9F_0$ at $d = 0.39h$, and we can assume there is no more attraction for distances $d > 2h$. In Table A.1 we report the value of F_0 we obtained for our prototype.

Note that these simplifications lead only to a small approximation error. Compared to an angle dependent formulation, a tilt of up to $\theta = 30^\circ$ leads to a max error in our model (Eq. A.14) equivalent to shifting the distance d by ± 3 [mm] (Figure A.1).

A.1.1.2 Angle-aware dipole-dipole model

In this section, we derive the complete EM model, using both, the pen position and its tilting angle as free variables. We continue the deduction of \mathbf{F}_p by substituting Eq. A.5—A.10 into the main expression Eq. A.4. However, by following that path we wouldn't necessarily attain information on how strong the actuation force depends on the tilting angles θ and φ . Here we take a different path. Based on the geometry of our system, we consider the cases where the pen is tilted by only a small angle. We introduce this small-angle approximation by keeping only the first order terms in θ ,

$$\sin \theta \approx \theta \quad (\text{with } \theta \text{ in radians}) \tag{A.16}$$

$$\cos \theta \approx 1 \tag{A.17}$$

As an indication of what this approximation means, for an angle $\theta = 30^\circ$, the difference between using $\sin \theta$ or $\cos \theta$ or their approximations forms (Eq. A.16 and

A.17) is 5% and 15 %, respectively. Under the small- θ approximation, the dipoles' vectors are,

$$\mathbf{m}_e = \alpha m_m \mathbf{e}_z \quad (\text{A.18})$$

$$\mathbf{m}_p \simeq -m_p \theta \cos \varphi \mathbf{e}_d + m_p \theta \sin \varphi \mathbf{e}_t + m_p \mathbf{e}_z \quad (\text{A.19})$$

and the distance between dipoles,

$$\mathbf{r}_p \simeq -(d + h_p \theta \cos \varphi) \mathbf{e}_d + h_p \theta \sin \varphi \mathbf{e}_t + h \mathbf{e}_z \quad (\text{A.20})$$

with the length of that distance, at first order on θ ,

$$r_p \simeq d^2 + h^2 + 2dh_p \theta \cos \varphi \quad (\text{A.21})$$

In turn, the scalar products (Eq. A.8—A.10) can be written as,

$$\langle \mathbf{m}_e, \mathbf{r}_p \rangle \simeq \alpha m_m h \quad (\text{A.22})$$

$$\langle \mathbf{m}_p, \mathbf{r}_p \rangle \simeq m_p [-\theta \cos \varphi d + h] \quad (\text{A.23})$$

$$\langle \mathbf{m}_e, \mathbf{m}_p \rangle \simeq \alpha m_m m_p \quad (\text{A.24})$$

We can now substitute these expressions into the main force equation A.4. As we do in previous section, we consider only the terms that contribute to the component \mathbf{e}_d of the force. Keeping only these terms that contain θ up to the first order,

$$\begin{aligned} \mathbf{F}_p^{(d)} &= \frac{3\mu_0 \alpha m_m m_p}{4\pi r_p^5} \left[-d + \frac{5dh^2}{r_p^2} - h\theta \cos \varphi - h_p \theta \cos \varphi + \right. \\ &\quad \left. + \frac{5h^2 h_p \theta \cos \varphi}{r_p^2} - \frac{5hd^2 \theta \cos \varphi}{r_p^2} \right] \mathbf{e}_d \quad (\text{A.25}) \end{aligned}$$

$$\begin{aligned} &= \frac{3\mu_0 \alpha m_m m_p}{4\pi (h^2 + d^2)^{5/2}} \left[\frac{-d(d^2 + h^2) + 5dh^2}{(h^2 + d^2)} + \right. \\ &\quad \left. \theta \cos \varphi \left(-h - h_p + \frac{5(h^2 h_p - hd^2)}{(h^2 + d^2)} - \frac{5d^2 h^2 h_p}{(h^2 + d^2)^2} \right) \right] \mathbf{e}_d \quad (\text{A.26}) \end{aligned}$$

$$\mathbf{F}_p^{(d)} = \alpha F_0 [f_0(d) + \theta \cos \varphi f_1(d)] \mathbf{e}_d \quad (\text{A.27})$$

where we define,

$$F_0 = \frac{3 \mu_0 m_p m_e m}{4 \pi h^4} . \quad (\text{A.28})$$

$$f_0(d) = \frac{d \left(4 - \frac{d^2}{h^2}\right)}{h \left(1 + \frac{d^2}{h^2}\right)^{\frac{7}{2}}} \quad (\text{A.29})$$

$$f_1(d) = \frac{1 + \frac{h_p}{h}}{\left(1 + \frac{d^2}{h^2}\right)^{\frac{5}{2}}} + \frac{5 \left(\frac{h_p}{h} + \frac{d^2}{h^2}\right)}{\left(1 + \frac{d^2}{h^2}\right)^{\frac{7}{2}}} - \frac{5 \left(\frac{h_p}{h}\right) \left(\frac{d^2}{h^2}\right)}{\left(1 + \frac{d^2}{h^2}\right)^{\frac{9}{2}}} \quad (\text{A.30})$$

Note that by considering the case $\theta = 0$ in Eq. A.27, we recover what we obtain for \mathbf{F}_a as calculated in Eq. A.14. That means that the equation for $\mathbf{F}_p^{(d)}$ subsumes the cases of the pen being tilted by a small angles, and it can be used in future EM actuated systems which may be able to track θ and φ .

B

MARLUI: MULTI-AGENT REINFORCEMENT LEARNING FOR ADAPTIVE UIS

B.1 LEARNED LOWER LEVEL

We detail how to learn the low-level motor policy. The low-level policy needs to learn i) the coordinates and dimensions of menu slots, ii) an optimal speed-accuracy trade-off given a target slot, and its current position.

To prevent the low-level motor control policy from correcting wrong high-level decisions and to increase general performance, we limit the state space S_M to strictly necessary elements with respect to the motor control task [47]:

$$S_M = (\mathbf{p}, \mathbf{t}), \quad (\text{B.1})$$

with the current position $\mathbf{p} \in I^2$, the target slot $\mathbf{t} \in \mathbb{Z}_2^{n_s}$.

The action space A_M is defined as follows:

$$A_M = (\mu_{\mathbf{p}}, \sigma_{\mathbf{p}}). \quad (\text{B.2})$$

It consists of $\mu_{\mathbf{p}} \in I^2$ and $\sigma_{\mathbf{p}} \in I$, i.e., the mean and standard deviation which describes the endpoint distribution in the unit interval. We scale the standard deviation linearly between a min and max value where the minimum value is the size of normalized pixel width and the max value is empirically chosen to be 15% of the screen width. Once an action is taken, we sample a new end-effector position from a normal distribution: $\mathbf{p} \sim \mathcal{N}(\mu_{\mathbf{p}}, \sigma_{\mathbf{p}})$.

Given the predicted actions, we compute the expected movement time via the WHO model [93], similar to our non-learned low-level motor control policy in the main paper.

The reward for the low-level motor control policy is based on the *motoric* speed-accuracy trade-off. Specifically, we penalize: i) missing the target supplied by the high-level ($-h$), and ii) the movement time (T_M). Furthermore, we add a penalty iii) which amounts to the squared Euclidean distance between the center of the target \mathbf{t} and $\mu_{\mathbf{p}}$. This incentivizes the policy to hit the desired target in the correct location. Since the penalty only considers the desired point $\mu_{\mathbf{p}}$, it will not impact the speed-accuracy trade-off (which is a function of $\sigma_{\mathbf{p}}$). The total reward is defined as follows:

$$R_M = \underbrace{\lambda(\neg h)}_{i)} - \underbrace{(1-\lambda)T_M}_{ii)} - \underbrace{\beta\|\mu_{\mathbf{p}} - \mu_{\mathbf{t}}\|_2^2}_{iii)}, \quad (\text{B.3})$$

where $\neg h$ equals 0 when the target button is hit and -1 on a miss. A hit occurs when the newly sampled user position \mathbf{p} is within the target \mathbf{t} , while a miss happens if the user position is outside of the target. λ is a speed-accuracy trade-off weight and β is a small scalar weight to help with learning.