Exploring the Making, Modifying, & Use of Physical Tools in Augmented Reality

Dominic Michael Potts, BSc (Hons)
School of Computing and Communications
Lancaster University

A thesis submitted for the degree of
Doctor of Philosophy

June, 2023

Declaration

I declare that the work presented in this thesis is, to the best of my knowledge and belief, original and my own work. The material has not been submitted, either in whole or in part, for a degree at this, or any other university. This thesis does not exceed the maximum permitted word length of 80,000 words including appendices and footnotes, but excluding the bibliography. A rough estimate of the word count is: 62663

Dominic Michael Potts

Exploring the Making, Modifying, & Use of Physical Tools in Augmented Reality

Dominic Michael Potts, BSc (Hons).
School of Computing and Communications, Lancaster University
A thesis submitted for the degree of *Doctor of Philosophy.* June, 2023

Abstract

The relationship between humans and physical tools is fundamental to all forms of work. With the proliferation of technology, the field of Human-Computer Interaction has focussed on the development of tools to interface with computers. Augmented Reality (AR) is one such technology that has gained prominence in recent years by merging physical and digital elements to enhance labour and promising more versatile and adaptive forms of work. While virtual objects provided by AR can enrich our physical environments, they remain intangible, presenting a number of challenges and opportunities for interaction.

Considering this, we first argue the necessity of adopting physical tools to interact with virtual objects in AR. Second, the design of physical tools in AR can yield entirely new interaction possibilities especially when combined with the physical environment. And third, we advocate for the use of physical but versatile tools in AR, capable of modification to meet the demands of the task and user.

This thesis explores *Tool-making*, *Tool-Modifying*, and *Tool-Using* for interacting with virtual objects in head-mounted display AR. The research includes a design space of physically-modifiable AR tools, supported by two empirical studies, and a toolkit for creating physically-modifiable cubic AR tools. We outline the surface-based touch gestures and 3D interaction techniques enabled by the cubic tools and combine them into five demonstrative AR applications. Lastly, we empirically explore and evaluate a 3D manipulation technique enabled by the toolkit.

Results of the design space exploration and studies showcase how physically-modifiable AR tools can engender novel forms of interaction, alleviate common limitations of physical interfaces, and address certain interaction challenges in current AR techniques. The aim of this thesis is to stimulate fresh considerations for AR design, bringing interaction to the physical environment where humans excel, and envisioning tools as the ultimate mediators in the ultimate display.

Publications

Contributing Publications

One publication, shown below, has been created directly from the thesis from which large portions of this published work is used within Chapter 4:

Dominic Potts, Martynas Dabravalskis, and Steven Houben. "TangibleTouch: A Toolkit for Designing Surface-Based Gestures for Tangible Interfaces". *Sixteenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '22. Daejeon, Republic of Korea: Association for Computing Machinery, 2022. ISBN: 9781450391474. DOI: 10.1145/3490149.3502263. URL: https://doi.org/10.1145/3490149.3502263

Related Publications

The following publications have been generated while developing this thesis and are partially related to the subject matter:

Dominic Potts et al. "ZenG: AR Neurofeedback for Meditative Mixed Reality". *Proceedings of the 2019 on Creativity and Cognition*. C&C '19. San Diego, CA, USA: Association for Computing Machinery, 2019, pp. 583–590. ISBN: 9781450359177. DOI: 10.1145/3325480.3326584. URL: https://doi.org/10.1145/3325480.3326584

Ludwig Sidenmark, Dominic Potts, Bill Bapisch, and Hans Gellersen. "Radi-Eye: Hands-Free Radial Interfaces for 3D Interaction Using Gaze-Activated Head-Crossing". *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. Yokohama, Japan: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445697. URL: https://doi.org/10.1145/3411764.3445697

Daniel Harris, Dominic Potts, and Steven Houben. "User-Elicited Surface and Motion Gestures for Object Manipulation in Mobile Augmented Reality". *Adjunct Publication of the 24th International Conference on Human-Computer Interaction with Mobile Devices and Services*. MobileHCI '22. Vancouver, BC, Canada: Association for Computing Machinery, 2022. ISBN: 9781450393416. DOI: 10.1145/3528575.3551443. URL: https://doi.org/10.1145/3528575.3551443

Thomas Wells, Dominic Potts, and Steven Houben. "A Study into the Effect of Mobile Device Configurations on Co-Located Collaboration Using AR". *Proc. ACM Hum.-Comput. Interact.* 6.MHCI (Sept. 2022). DOI: 10.1145/3546735. URL: https://doi.org/10.1145/3546735

Unrelated Publications

The following publications have been generated while developing this thesis. While unrelated to the topic these works to an extent have guided the thesis into what it has become:

Maria L. Montoya Freire, Dominic Potts, Niraj Ramesh Dayama, Antti Oulasvirta, and Mario Di Francesco. "Foraging-based optimization of pervasive displays". *Pervasive and Mobile Computing* 55 (2019), pp. 45–58. ISSN: 1574-1192. DOI: https://doi.org/10.1016/j.pmcj.2019.02.008. URL: https://www.sciencedirect.com/science/article/pii/S1574119218307570

Kim Sauvé, Dominic Potts, Jason Alexander, and Steven Houben. "A Change of Perspective: How User Orientation Influences the Perception of Physicalizations". *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20. Honolulu, HI, USA: Association for Computing Machinery, 2020, pp. 1–12. ISBN: 9781450367080. DOI: 10.1145/3313831.3376312. URL: https://doi.org/10.1145/3313831.3376312

Claudia Daudén Roquet, Corina Sas, and Dominic Potts. "Exploring Anima: a brain–computer interface for peripheral materialization of mindfulness states during mandala coloring". *Human–Computer Interaction* 0.0 (2021), pp. 1–41. DOI: 10.1080/07370024.2021.1968864. eprint: https://doi.org/10.1080/07370024.2021.1968864. URL: https://doi.org/10.1080/07370024.2021.1968864

Edward Thompson, Dominic Potts, John Hardy, Barry Porter, and Steven Houben. "AmbiDots: An Ambient Interface to Mediate Casual Social Settings through Peripheral Interaction". *Proceedings of the 33rd Australian Conference on Human-Computer Interaction*. OzCHI '21. Melbourne, VIC, Australia: Association for Computing Machinery, 2022, pp. 99–110. ISBN: 9781450395984. DOI: 10.1145/3520495.3520504. URL: https://doi.org/10.1145/3520495.3520504

Acknowledgements

A number of people have been ‘instrumental’ in the development of this thesis. I am immensely grateful to my supervisors, Dr. Steven Houben and Prof. Hans Gellersen, whose unwavering support and profound wisdom have made this thesis possible. Collaborating with them throughout my PhD journey has been an absolute privilege, and their guidance has shaped me as a researcher. A sincere thank you for nurturing my passion for the field of Human-Computer Interaction, pushing me to new limits, and imparting invaluable skills and knowledge that have enriched my academic journey.

My exceptional colleagues and friends in academia, both past and present, have played a vital role in elevating my work and ideas. Their presence has fostered a sense of camaraderie and community throughout my PhD journey, providing me with the necessary space and support to evolve as a researcher, a teacher, and an individual. I would like to extend my gratitude to Chris, Ludwig, Argenis, Pierre, Rob, Tom, Claudia, Josh, Franzi, James, and others for generously lending their ears, minds, and hearts. I would also like to express my appreciation to those outside the academic sphere who have been integral to my life and have shared in both the ups and downs. Jeremy, Alex, Edward, Eddie, and Ben, I am privileged to count you among my friends.

This thesis is dedicated to my family, whose love and support have made my PhD journey possible. To my Mum, Dad, brother Jonny, and Grandad, you have been my steadfast advocates from the beginning, providing me with the care, resources, and love to achieve this milestone. I would also like to extend my gratitude to my wider family, particularly my Uncle Bobby and Rob, whose influence and encouragement during my formative years instilled in me a curiosity for the world and a deep appreciation for academia. Above all, to my wonderful partner Kim who has been a substantial and unshakeable driving force behind me every step of the way through her love, support, and laughter.

Finally, during the course of my PhD, our family experienced the loss of my Aunty Mavis after her battle with cancer during the COVID-19 pandemic. Her indomitable spirit through the toughest of times is something I can only aspire to emulate. This thesis, with deep reverence and love, is also dedicated to her.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This work is framed by exploring the concepts of physical tool *making*, *modifying*, and *use* in Augmented Reality (AR). Our relationship with physical objects and tools has been of profound historical importance in facilitating physical labour resulting in large technical, cultural, and societal shifts in human development. This relationship has endured and evolved in modern times shaping how we operate in both domestic and work environments as the tools at our disposal become increasingly sophisticated.

Particularly, the emergence of digital technology is one of the greatest examples in recent memory of how our relationship with tools has evolved as devices become increasingly ubiquitous, embedded, and all-encompassing in our daily lives. Human-computer interaction (HCI) as a research field emanated from this burgeoning and proliferation of technology, concerning itself with the development of tools to interface with machines in our environment, ultimately expanding human capabilities. This is to say that tool development is a fundamental component of HCI.

Technological developments in recent decades have led to the emergence of Augmented Reality (AR), a technology that blends physical and virtual elements, offering new possibilities for work and recreation. With the improvement of AR display and tracking capabilities the field of HCI has seen a growth of research opportunities, particularly in AR interaction design. To expand user functionality in AR and virtual environments, multiple approaches to interaction have been developed, with a subset of research focussing on using physical devices for interaction — *physical tools for virtual spaces.*

The combination of AR environments and physical tools not only enables direct physical interaction with virtual objects but also has the potential to transform our relationship with tools and how we approach digital and virtual work. An emerging possibility, and a central theme of this thesis, is the concept of *virtually* and *physically-modifiable* tools. Virtual modification refers to the remapping of a physical tool's virtual function and abilities, while physical modification refers to the alteration of a physical tool's form factor as a mechanism for input.

This thesis makes a case for the adoption of **physical** tools to facilitate interaction in AR as many others have before [46, 191, 194]. More specifically, the work delves into Napier's concepts of *Tool-Use*, *Tool-Modifying*, and *Tool-Making* [275], for labour not in physical environments but for virtual environments such as AR. In addition, the thesis explores tools that are both *physically* and *virtually modifiable* to engender both 'versatile' and 'real' interaction [194] and operate as a mediator where physical and virtual elements coalesce such as AR environments.

1.1 The Past, Present, and Future of Tools

Take a moment to admire the human hands. The structure of the hand has been moulded by our relationship with the physical environment for millions of years since it is theorised that our arboreal ancestors became terrestrial [154]. As John Napier describes, hand *structure* dictates *function* — 'the integrated action of physical parts' — and hand function dictates how we enact on the *environment* [275]. During the early stages of primate evolution, the development of prehensile patterns allowed for an extension of function in the environment. In humans, this evolution reached a milestone with the prevalence of dual prehensile patterns. While the ***precision*** and ***power*** grip can be seen as basic ways of categorising prehensile movements, it is important to note that the hands are capable of a wide range of actions. But it is not just the hands alone that extend function. The combination of the hands and the physical objects in our modern and ancient environments leads to a diversity of movements and the eventual adoption, modification, and creation of *tools*.

Starting as analogous materials shaped by requirements and imagination, *tool-use* became inherent to all labour and later became deeply intertwined with *tool-making*. As a critically definitive human characteristic and hallmark of collaboration, *tool-making* has been maintained between neighbours and generations by explanation, example, and demonstration. Our relationship with tools is as old as the human experience and something that has been sustained in modern times. Labour has been continuously accommodated as we develop increasingly complex tools, shaping the environments where we work and recreate.

Digital technology and its proliferation have brought about a revolutionary change in the way we work and have greatly expanded the definition of tools. By understanding the phenomena around people and technology, HCI has helped to create digital tools we now find commonplace, from the inception of the modern work station [109, 349], to GUI and WIMP desktop interfaces, and the prevalence of mobile computing. The prevalence of digital tools and interfaces has resulted in somewhat of a renaissance in virtual *tool-making* with increasing accessibility and capabilities while relying on a standardised set of physical tools — such as the keyboard, mouse, and mobile devices.

One promising technology that has the potential to expand human *function* by overlaying virtual information on physical environments is Augmented Reality (AR). HCI researchers have amassed a wealth of knowledge about interface design for traditional desktops, which is highly applicable to the design of 3D user interfaces and AR [53, 278, 281, 342]. The research on 3D UI has played a crucial role in the development of AR interaction [221], particularly in creating gestural and controller-based interfaces that seamlessly blend real-world and new metaphors. Significant progress had also been made in the 1990s on tracking and display technology for AR [41, 48, 323, 371]. However, the field of AR has focused mainly on ”purely visual augmentations” [191] and is still heavily dependent on visual perception [227]. As a result, interaction with AR environments is still in its infancy and is mainly limited to viewing or browsing virtual information overlaid on the real world. As physical and virtual spaces become more integrated, interacting with virtual content and objects in AR will become increasingly essential, similar to how interacting with WIMP interfaces became crucial for desktop work.

This realization about virtual environments has led to the development of new interaction models such as ‘Instrumental Interaction’ [31], ‘Reality-based interaction’ [194], and Mixed Interaction [91]— a call for *tool-making* in a post-WIMP era. These new interaction models are based on laws, rules, and conventions from the physical world, allowing users to apply their understanding of the real world when interacting in virtual environments. This effectively reduces cognitive or functional barriers between a user and their goals, also known as the “gulf of evaluation and execution” [185]. More specifically, these models leverage our understanding of our bodies, spatial reasoning and proprioception, our relationship with physical space and other people, and our natural tendency to use physical tools. The advent of Tangible User Interfaces (TUIs) is an ideal example of exploiting “humans’ most basic knowledge about the behaviour of the physical world” [389] to provide critical sensory cues for interaction. As a result, many have made the case for adopting physical tools (Tangibles) to mediate the virtual and physical elements of a blended AR environment and remedy the existing interaction challenges [46].

In some ways, physical tools have already been adopted in commercial AR systems in the form of controllers, but using a single tool with a fixed form factor for multiple functions has limitations. Tools are not necessarily a “one-size fits all” solution. Current interaction techniques for AR systems can be compared to the flaked stone tools and hand axes of early humans, only allowing basic forms of interaction with the environment to facilitate simple labour. Despite the advancements in Tangible AR, it remains unclear how future physical tools should be designed, modified, and operated to facilitate more advanced forms of labour in virtual environments such as AR. As a result, the precursor to *tool-use* in AR should be *tool-making* and *tool-*

modifying. There are several conceptual and technical challenges involved in tool-making, including determining the various 'jobs' that AR tools should be able to perform, designing their affordances, and figuring out how to fabricate and track them in 3D space. Additionally, Jacob et al. highlights the inherent trade-off between 'versatile' digital interfaces (GUIs) and 'reality' interfaces (TUIs) [194]. However, both versatility and reality can be achieved if AR tools are designed to be both *physically* and *virtually modifiable* from the outset.

1.2 Research Aim and Questions

The primary working assumption in this thesis is that interaction with virtual objects in Virtual Environments will be *frequent* and *essential*. Interactions with virtual objects should match or exceed the precision and expressiveness we can achieve in the physical environment. In light of this assumption, the position presented in this thesis is that the creation, modification, and adoption of **physical tools** is not only *vital for addressing existing interaction challenges* in AR but will also become *salient in the future of labour*. To support this argument, we will first build a case for physical tools in AR by outlining three challenges that current AR interaction faces.

Fatigue : Current AR interaction is not sustainable due to fatigue. So-called '*gorilla-arm*' [188] is a negative effect related to arm fatigue during prolonged use of a mid-air interface. The '*consumed endurance*' model[164] describes how users can hold their arm fully extended at 90° for a little over 1 minute before inducing fatigue, and likewise at 45° for a little over 2 minutes. This is particularly relevant for AR due to popular interaction techniques relying on mid-air semantic and deictic gestures [200, 252], such as in Microsoft's Hololens [266], for direct manipulation. These interaction techniques are suitable for short bursts of interaction, but for any prolonged activities such as 3D modelling and manipulation of virtual objects, these techniques quickly become unsuitable due to lack of precision and physical fatigue [25, 195]. While a number of techniques and design guidelines have been proposed to reduce fatigue when using mid-air interfaces [150, 164], there appears to be a tradeoff between fatigue-inducing, direct gestures and more sustainable indirect gestures, such as micro-gestures [145], where direct hand-to-display coupling becomes lost. Alternatively, some AR systems rely on bespoke controllers for interaction, such as the Magic Leap [246]. However, these devices are also designed for mid-air interaction resulting in the same fatigue issues as mid-air gestures. If AR is to become truly ubiquitous and facilitate the future of work, supporting prolonged and direct interaction with virtual content is vital.

Combined sensory cues : When acting in the physical world, there are numerous perceptual cues, such as visual, auditory, tactile, proprioceptive, and vestibular cues, that aid our understanding of environmental constraints, affordances for action, and spatial reasoning. One important aspect of sensory perception is the manner in which different cues are *combined* and *integrated* to provide a more complete and precise picture of the environment [113]. Pertaining to virtual environments, researchers have been aware of inherent challenges people experience when understanding and performing actions in virtual 3D spaces [158]. This is a consequence of current AR/VR technology being unable to accommodate all perceptual cues that are important for different types of sensory 'disambiguation' and, as a result, 3D interaction is difficult to the extent that regular WIMP interfaces are insufficient [221]. In some aspects, visual and auditory sensory cues have been implemented to convey depth in AR/VR environments such as visual disparity, shadows, shading, motion parallax, occlusion, and spatialised audio. However, prior work has found that AR/VR applications are overly reliant on visual perception [227] when in actuality both visual and haptic information is leveraged for spatial interactions [112], with Mechanoreceptive feedback as a key aspect of precise motor-control [198]. To cater for combined sensory cues a number of solutions have been proposed from haptic wearables [94, 207], to instrumented environments [67, 94], to tangible artefacts and tools [46] all of which have various compromises and tradeoffs.

The Physical Environment : Current virtual environments tend to treat the physical environment as an obstacle, while AR has a unique ability to leverage it to support interaction. There are several examples of "*Opportunistic Controls*" being developed for AR [105, 106, 156, 157, 416], allowing users to gesture on and receive feedback from physical objects in their immediate environment. In this way, unused affordances can be leveraged to provide passive haptics to ease or enhance gestural input. However, the controls at the user's disposal are context dependent and also assume an awareness and tracking of the physical environment. Additionally, there also views of AR to be *nomadic* across physical environments [143, 213], in which case the environment might not always provide the same controls or input elements. AR tools can be anchored in the real world while also providing a way to interact with the virtual environment, allowing users to take advantage of the benefits of both physical and virtual spaces. As AR becomes more commonplace, there may be other people present in our physical spaces who do not have the same view of the virtual environment. This can be influenced by the proliferation of display technology and the ability of others to see virtual content. The use of AR tools not only has the capacity to mediate the physical and the virtual, but they can facilitate collaboration between users even if their views of the virtual environment are different. These tools can provide a shared reference point, or 'handle', that persists in the physical world.

This thesis argues, as many others have [46, 191, 194], that current interaction challenges in AR can be addressed by the adoption of physical tools. Considering these three interaction challenges, the overarching research aim is to *explore physical tools for AR* through concepts defined by Napier: *tool-making*, *tool-modifying*, and *tool-using* [275] with a focus on bridging the physical and virtual to facilitate labour in AR. Through this work, we aim to gain a better understanding of how these concepts should operate in virtual environments as opposed to physical ones.

- **Tool-Using**: "*An act of improvisation in which a naturally occurring object is utilised for an immediate purpose*". Just as tools expand our function in the physical world they have the same capacity in virtual environments. For AR, 'naturally occurring objects' can be 'everyday' objects that inhabit work and domestic spaces, bespoke instrumented artefacts, or a combination of the two.

- **Tool-Modifying**: "*Adapting a naturally occurring object by simple means to improve its performance*". In AR, a tool has the capability to be modified physically and virtually to not only improve performance in a given activity but also be multiplexed for a large number of tasks. Physical modification describes the adaptation of a tool's form factor, which can be achieved via shape-changing mechanisms or through a modular form factor. Virtual modification describes the adaptation of virtual affordances granted to a physical tool as well as the relationship between the input and output mechanisms. In both cases, Tool-Modification in AR should support ad-hoc repurposing of a physical object.

- **Tool-Making**: "*An activity by which a naturally occurring object is transformed in a set and regular manner into an appropriate tool for a definite purpose.*". From ancient humans in the throes of constructing rudimentary stone tools to refined personal creations of craftsmen, tool-making spawns from a combination of the imagination and an analysis of the activity and conditions the tool will be subject to. A '*set and regular manner*' or process should exist for Tool-Making in AR, as well as an understanding of what '*definite purpose*' a tool should have based on its, inherent or designed, affordances and properties.

Inspired by these concepts introduced by Napier [275], the work presented in this thesis asks and aims to answer three research questions around tool-using, tool-modifying, and tool-making for AR.

RQ1: *How do we ideate tools for Augmented Reality that are both physically and virtually modifiable?* ***Tool-making*** — The creation of tools has been a long-standing practice in the physical world, and in the digital realm, HCI has produced a variety of interaction devices, gestures, and techniques using various methodologies [222]. Similarly, there has been a significant amount of research on the development of tools for use in AR [43, 209]. However, our focus is on the creation of modifiable tools for AR. Previous work on shape-changing interfaces has shown potential for physically adaptable interfaces [321], and there are many examples of everyday objects that can be modified or shape-changed. To achieve this, we plan to explore methods for creating physical tools for AR by using existing physical objects as a way to probe the usefulness of certain properties and affordances for different tasks.

RQ2: *How should physically and virtually modifiable tools be created and operated for Augmented Reality?* ***Tool-Modifying*** — In this question there are several technical challenges to overcome: the fabrication, interaction detection, and 3D tracking of modifiable tools for interaction in AR. Additionally, such tools yield a rich design space for their operation facilitating new interaction techniques, metaphors, and applications for AR.

RQ3: *How do newly designed tools compare to existing interaction techniques and methods for Augmented Reality?* ***Tool-Using*** — There are a number of interaction challenges for virtual environments and AR specifically which are currently being researched [134, 420]. Knowing the benefits of physical interaction devices we explore: how these newly-designed tools afford new types of interaction, how they compare to state-of-the-art interaction techniques, and how they help to address the existing challenges for interaction.

1.3 Methodology

The overall goals of this research are to design, develop, and evaluate AR tools, interaction techniques, and applications, with a focus on exploring their potential and utility. The research questions are primarily exploratory in nature and are intended to help us better understand how AR tools can be created and used effectively in AR environments. As an overview, the work presented in this thesis involves developing methods for designing and prototyping AR tools, evaluating their usefulness through demonstrations [226, 298], and empirically evaluating proposed interaction techniques using these tools. We will primarily be using head-mounted display (HMD) AR as a platform for exploring tool-based interaction in virtual environments, as HMDs allow users to have their hands free while in the virtual environment, which is necessary for interacting with physical objects. HMDs also offer high-fidelity displays and tracking capabilities and provide full control over the experimental environment and conditions for empirical evaluations. While we use this specific platform, the findings and contributions within each section of work may have implications for the broader field of Mixed Reality [354].

Starting with **RQ1** (*Tool-Making*), the focus of Chapter 3, we begin by conducting a structured design space exploration of existing, shape-changing physical objects to inform the design of AR tools. We refer to these as physically-modifiable objects (PM-Objects) and we used a set of 20 to explore the design space of physically-modifiable tools (PM-Tools) for AR. The PM-Objects themselves ranged from sensory toys to gadgets and fidgets. We first produced a framework and vocabulary to describe and classify the PM-Objects, which can also be generalised to any type of physical object. We introduce the concepts of *augmenting* and *reconfiguring* PM-Objects, each embodying different affordances and properties which can potentially be leveraged for different types of AR tasks.

We then refined the set of 20 PM-Objects to 10, to cover all the different *augmenting* and *reconfiguring* properties and affordances with as few objects as possible. We then use these 10 PM-Objects to probe the mapping of object affordances to AR tasks, specifically 3D object manipulation and modelling. We conducted 4 group study sessions inspired by the Research through Design (RtD) method of 'card-sorting' [295] and guessability methodology [308, 418, 419] with a total of 10 participants. Participants were shown examples of the AR interactions on a display and then asked to collectively explore the PM-Objects, determine how they would perform the task with each object, and then rank the objects from the worst to best for that task.

Based on the results of participants' ranking of the PM-Objects, we conducted a guessability study inside of HMD AR with individual users using PM-Objects from the workshop that was ranked favourably. Using user-elicitation [419], we produced a consensus set of gestures and interactions using the objects, a total of 25 unique gestures, which we then consolidated into two conceptual examples of AR tools. Our consensus set and AR tool examples demonstrate how *augmenting* and *reconfiguring* object properties can be applied differently in the design of PM-Tools for different types of AR tasks.

For **RQ2** (Tool-Modifying) we focus on exploring the design space of one specific form-factor to mediate interaction in AR — cubes. In doing so we can explore the fabrication and prototyping [32] of *physically* and *virtually modifiable tools* in AR through a form factor that is conceptually, technically, and structurally as simple as a cube. **RQ2** is informed by the outcomes of **RQ1** but is steered more to the technical design, implementation, and deployment of different interaction artifacts [114]. For our work, the artifacts include a prototyping and fabrication toolkit, example interactive cubic tools, interaction techniques using the cubic tools, and applications [186]. The toolkit itself provides novel fabrication and rapid prototyping methods for designing modular, interactive cubic tools with interchangeable capacitive faces that facilitate surface-based gestures. We then evaluate the toolkit by demonstration [186] using several example interaction techniques and applications both inside and outside of AR. Prototyping different examples allows the exploration of different interactions and ultimately allows the design space of *modifiable tools* to be broadened and refined.

Finally, **RQ3** involves a set of laboratory user studies [102] exploring and evaluating a precise 3D object manipulation technique using the cubic tools produced for **RQ2**, which is described in Chapter 5. The interaction technique itself is an asymmetric bimanual technique that uses two cubic tools, one operated in the user's hand and the other operated in conjunction with a physical surface. The cube in hand operates as a direct spatial proxy to a virtual object (3D translation and rotation) while the other acts as a 'dial' to adjust the control-display gain, giving users explicit control over the precision of a manipulation. Both cubes can also be used together to uniformly and precisely scale a virtual object, again using control-display gain.

We conducted three studies to 1) compare the efficiency and user preference of three variations of the technique design, 2) explore unique behaviour that arises when using the technique in terms of participant handedness, and 3) comparatively evaluate the new interaction technique against two baseline techniques that use AR/VR hand-controllers. For each study, we recorded a mixture of quantifiable performance measures (such as time and accuracy) and self-reported measures (such as subjective user preference/workload scores) and combined these with qualitative participant comments and researcher observations. The goal of these studies was to provide insight into new behavioural phenomena surrounding the interaction technique and ultimately benchmark it against state-of-the-art AR techniques.

1.4 Contribution

The work in this thesis presents the following contributions:

- **Framework and Methodology for creating Physically Modifiable Tools for AR interaction (RQ1 & RQ2).** We introduce and explore the concept of *physically-modifiable objects*, shape-changing objects leveraged as interaction tools for complex tasks in head-mounted AR, such as 3D manipulation and modelling. We propose new vocabulary for systematically classifying existing physical objects, with modifiable or shape-changing properties, and explore how their properties and affordances can be mapped to complex AR tasks. Using a representative set of 10 exemplar physical objects, we conducted a group-design study (*Study 1*) and a follow-up AR user-elicitation study (*Study 2*) to map object affordance to AR tasks. User gestures and actions with tangible objects are classified into a new elicitation taxonomy and then into a consensus set. Finally, we operationalise our consensus set and design considerations in a set of conceptual PM-Tools for virtual object-centric interactions in AR. The insights from the studies and example PM-Tools provide a deeper understanding of how shape-change and physical modification can be leveraged for input and how different properties and affordances can be appropriated into new form factors for interaction.

- **A tangible toolkit for fabricating and prototyping modifiable cubic artifacts for AR interaction (RQ2 & RQ3).** We contribute a novel artifact — a rapid prototyping and fabrication toolkit for modifiable, modular, interactive cubes that have interchangeable capacitive faces designed to facilitate a multitude of surface-based gestures. We then adapted the toolkit to provide a novel method for tracking physical objects in 3D, providing the basis for exploring physically and virtually modifiable cubic tools in AR.
  - **TangibleTouch Toolkit:** We present a modular and extendable hardware-software platform enabling the rapid fabrication of cubic tangibles capable of detecting different surface-based gestures. We describe the design of the cubes and interchangeable faces, the hardware and fabrication process, the configuration of surface gestures, and the deployment of gestures as interaction techniques. We evaluate the *physical-modification* capabilities of the cubic artifacts through three small demonstrative applications, displaying the range of surface-based gestures supported and highlighting the generality of the toolkit (**RQ2**).
  - **3D AR Tracking:** We extended the toolkit to incorporate a fusion of three different tracking approaches, computer vision, IMU, and AR hand-tracking, in order to appropriately track interactive cubes in 3D and explore them for interaction. We then characterise the precision of this tracking approach (**RQ3**).

- **Design Space and Applications of using Cubic Tools in AR (RQ2 & RQ3).** We contribute a design space exploration of AR interaction with a set of identical physical artifacts that are cognitively and structurally as simple as a cube. We then demonstrate the breadth of interactions cubes enable in AR as well as their *virtual-modification* capabilities through two demonstrative AR applications — *AR Workspace* and *AR Maps*.

- **Design, demonstration, and empirical evaluation of a technique for AR manipulation using cubes(RQ3).** Using the novel toolkit we developed, we produced an interaction technique for performing precise 3D virtual object manipulations in AR. We present an exploration of the technique through an empirical study of different technique designs (*Study 1*) and a behavioural study around user handedness when performing manipulations (*Study 2*). We then conduct a comparative evaluation (*Study 3*) of the cube interaction technique against a set of state-of-the-art controller-based techniques.

1.5 Thesis Structure

This thesis follows the structure below:

Chapter 2 provides an in-depth examination of the field of Human-Computer Interaction, with a focus on Augmented Reality (AR) and Tangible Interaction. The chapter starts by giving an overview of the state of the art in AR research within the context of HCI, highlighting key interaction techniques and current challenges. We delve into two particularly important challenges: the need to account for multiple sensory cues and the challenge of reducing complexity in 3D interaction. Next, the chapter explores emerging interaction models that can be applied to AR and summarizes recent advancements in Tangible User Interface (TUI) research as a foundation for exploring physical and modifiable tools in AR. We conclude by discussing the technical challenges and existing approaches for creating physical tools for AR. This chapter serves as a comprehensive introduction to the current state of AR and TUI research, providing a solid foundation for the rest of the thesis.

Chapter 3 outlines the design space and framework for exploring physically-modifiable tools (PM-Tools) for AR interaction. It explains the process for exploring the design space using existing physically-modifiable objects (PM-Objects) and evaluates their potential and capacity to be used as input devices for AR activities. The chapter then delves into the two empirical studies that form the basis of the design exploration of PM-Tools, summarizes the outcomes, and synthesizes the findings into conceptual examples of PM-Tools.

Chapter 4 introduces the design space of a specific tool form factor for AR — cubes. We identify three interaction metaphors that can be applied to virtually modify the tools and also present the broad range of surface-based gestures that are supported by the cubic form factor. From this, we detail our rapid fabrication and prototyping toolkit (*TangibleTouch*) that produces physically-modifiable cubic tools that are capable of detecting surface gestures and being tracked in 3D. Throughout the chapter, we showcase the interactions and applications afforded by the cubic tools and finish by reflecting on the toolkit and opportunities for it to be expanded. Part of the work presented in this chapter was originally published in the Sixteenth International Conference on Tangible, Embedded, and Embodied Interaction (**TEI'22**).

Chapter 5 details three empirical lab studies that are the basis of an in-depth exploration and evaluation of an asymmetric bimanual technique for precise 3D virtual object manipulations using cubic tools. The chapter begins by outlining the implementation and operation of the technique and then provides an in-depth analysis of the first study that investigates the impact of three different control-display gain designs. The chapter continues by describing the second study that delves into the behavioural phenomena around user handedness when operating the technique. The chapter concludes by describing and presenting the results of the third study that evaluates the new technique against established AR/VR hand-controller techniques.

Chapter 6 discusses the work presented in the previous chapters and reflects on the research questions. It provides details on the outcomes of the thesis and specifically what was learned about the making, modifying, and using of physical tools in AR. We also highlight the limitations of our work and present opportunities for future work on physical AR tools.

Finally, we conclude the thesis in **Chapter 7**.

Chapter 2

Background

This thesis investigates the use of *physical tools* in Augmented Reality (AR) from the perspective of HCI. First, the work is contextualised within the field of head-mounted display AR by defining key terms and providing an overview of current AR interactions, challenges, and research. The concept of 'tools' in HCI, Tangible User Interface research, and the integration of Tangibles with AR are then examined. The chapter continues by reviewing current research on shape-changing interfaces as the foundation for *physically modifiable* tools. Finally, the chapter concludes by discussing the current approaches and challenges to developing physical tools for AR, including prototyping, sensing, and tracking.

2.1 Defining AR, VR, VE's and MR

A first glimpse of the possibilities of Augmented Reality came when Ivan Sutherland speculated in the late 1960s on the concept of an '*Ultimate Display*' [368]. Sutherland first describes our relationship with the physical world, the properties of which are predictable and familiar, and ruminated on the possibilities of future '*kinesthetic*' displays. Interfaces that go beyond crunching numbers, instead engendering interactive experiences to make what is typically intangible **tangible**. A long evolution of Augmented Reality was kick-started by a number of pioneering systems [52], one of which was developed by Sutherland. The '*Sword of Damocles*' [369] head-mounted display (HMD) combined with head tracking allowed the user to look around a simple virtual world of primitive wireframe objects. In actuality, the display yielded two different ideas. If operated in low-light conditions the user would only see the immersive virtual world, what we now describe as Virtual Reality (VR), and if operated in normal lighting the user would see the physical world overlaid with the virtual world, now described as Augmented Reality (AR) [30].

In the 90s, Milgram took this further envisioning these two concepts on a spectrum often referred to as the Reality-Virtuality continuum [269, 270], with AR encompassing the area between **real**, physical environments and entirely simulated, **virtual** environments. While modern understanding still supports these definitions of AR/VR, there have been more recent attempts to update and consolidate the terminology of the field [354] in light of more recent technological advancements and to encompass more than just *visual* displays. For consistency and posterity, these are the current terms and definitions that contextualises the work of this thesis and are used throughout.

Virtual Reality (VR) : Fully '*synthetic view*' of an entirely '*constructed*' world or reality [354]. VR is typically characterised by the use of head-worn displays with head tracking that are entirely immersive. In relation to Milgram's continuum, this is the far end of the reality-virtuality continuum [270]. VR is probably the most agreed upon and understood definition in the field.

Augmented Reality (AR) : Primarily defined as the '*Merging of 3D graphics with the real world*' [354] or the '*seamless blending of virtual images with the real world*' [43]. Particularly with AR, there is an emphasis on '*seeing*' the real world and unlike VR must always '*happen in the physical space*' [23, 43, 270, 354]. These fundamentals are a point of departure as there is ongoing discussion and debate within the field about what constitutes AR and what mediums and experiences fall under the umbrella. For example, some argue spatial registration of virtual objects is not always necessary so long as overlays are 'contextual' [354]. Additionally, some argue that interaction with the physical space is essential for both users and virtual content while using AR. Regarding display technology, there is a multitude of mediums to experience elements of AR. For brevity, we focus on the archetypal visual mediums which range from Head-mounted Displays (HMD) to projection-based systems (sometimes referred to as Spatial AR [49]), to Mobile AR. As described in Section 1.3, the thesis will focus on HMD AR that allows for both '*seeing*' and interacting with the physical world.

Virtual Environments (VEs) : An interactive, computer-generated simulation of three-dimensional spaces. Historically, the term was synonymous with VR or strongly implied the utilisation of head-tracking and head-worn displays [23, 108, 347]. However, in this thesis, we use the term to encompass both VR and AR environments that are accessed through head-mounted displays and feature spatial tracking of the user. This broad definition allows for a more inclusive discussion as there may be relevancies for both VR and AR environments.

Mixed Reality (MR) : Milgram used MR as an umbrella term to encompass the spectrum between entirely virtual and physical environments [270]. However, in recent years there has been significant debate surrounding the definition and understanding of Mixed Reality in the field of HCI [354]. Some experts argue that MR is an updated understanding and description of AR, with a focus on the interaction between virtual and physical elements, as opposed to AR being a purely visual overlay that is detached from the physical environment [325, 354]. Others believe that true MR should allow for fully immersive VR and AR experiences, as well as the transitory state between them [297]. There are also those in the field who use the term MR synonymously with AR [354]. Due to this ongoing ambiguity in definition, we avoid using the term *Mixed Reality* in this thesis for the most part.

2.2 AR/VR Technology and Current Applications

A myriad of different AR technology has been developed from large immersive CAVE (Cave Automatic Virtual Environment) systems and projection-based AR [49, 52, 92, 293], to head-mounted optical and video-see through displays [23, 65], to reduced form factors such as smart-glasses [97, 208] and even mobile device AR [74, 152, 410]. Nearly 60 years after Sutherland's formative prototype [369], AR and VR technology using head tracking has become increasingly refined and commercially accessible in the form of consumer AR (e.g. MagicLeap [1] and Hololens [2]) and VR HMDs (e.g. Vive [3], Quest [4] and Index [5]). Many AR devices feature full environmental simultaneous localization and mapping (SLAM) tracking [374], along with the option to incorporate marker-based tracking [121] without the need for external sensors. These headsets also support various forms of interaction, including hand gestures, gaze tracking, and voice commands, with some even providing specialised handheld controllers for interacting with virtual content. Given that these headsets offer the highest fidelity AR experiences and offer robust head tracking, the research in this thesis covered in Chapters 3, 4, and 5 is conducted using AR HMDs.

In terms of applications, VR has found a niche in immersive entertainment and training/simulation [52, 280] seeing increasing adoption both domestically and industrially. There are also many promising applications for AR. In the healthcare industry, AR can be a hands-free, assistive interface during complex procedures or surgery [253, 433]. In manufacturing, AR can be used to provide workers with visual instructions or be used to mediate co-presence with a remote expert [257, 280]. In education, similar to VR, AR can be used to enhance learning by providing students

[1] MagicLeap 2: https://www.magicleap.com/magic-leap-2
[2] Microsoft Hololens 2: https://www.microsoft.com/en-us/hololens
[3] HTC Vive: https://www.vive.com/uk/
[4] Meta Quest: https://www.meta.com/gb/quest/quest-pro/
[5] Valve Index: https://store.steampowered.com/sub/354231/

with interactive and immersive experiences that, unlike VR, can be grounded in the physical teaching space [42, 77, 420, 232].

Despite advancements in display quality and capabilities, AR has encountered more obstacles to widespread adoption compared to VR [6, 9, 100, 134, 333]. Additionally, even though both AR and VR have been present in the public consciousness for decades, the integration of AR into everyday work practices has not been as successful as the adoption of personal computing, smartphones, and the Internet of Things. There are a number of factors impacting AR adoption but a primary factor, and a focus of this thesis is the challenges users encounter when interacting in AR. For all the strides made in display capabilities and tracking, interaction for both AR and VR remains somewhat in infancy.

2.3 AR Interaction

First, let us be clear with what we mean by '*AR Interaction*'. Interaction as a concept in HCI is a broad topic of ongoing discussion [177] and there are many interpretations that are foundational to key HCI theories. Hornbæk et al. provides an extrapolation of different understandings of interaction from existing concepts [177] and how these can lead to different interpretations and measures of '*good*' interaction. For the purpose of this thesis, we will mainly focus on three interpretations of interaction — interaction as *dialogue* [64, 285], *information transmission* [127, 245], and most importantly *tool-use* [31, 53, 218, 282].

Interaction as dialogue describes the 'cycle of communication actions channeled through Input and Output' [177]. This understanding is prominent and historical in HCI, viewing interaction as a set of stages or turn-taking between the human and the computer based on a user's formulated goals. This interpretation is prevalent in GOMS [64] and Norman's action model [285]. These models stress the importance of user actions to be understood by the computer and vice versa, something Norman describes as the 'gulf of evaluation and execution' [284]. As such, interface *mapping* is a crucial concept — the ability for a user to understand how to achieve their goals using the interface. We expand on mapping in Section 2.5.1. In this case, *direct* mappings are preferable to provide a 'strong sense of understanding and control' [185, 341]. In the case of AR interaction, guessability studies [419] and user elicitation [308] are commonly employed methods to evaluate and design 'good' interaction, which is typically comprised of more *direct* mapping. Ultimately, the activities the user engages in (something we expand on in Section 2.3.1) and their goals whilst in AR determines how an interface should be designed and constructed so that the mapping maintains *directness*.

Interaction as information transmission describes the relationship between the user and computer not as a series of stages but rather the rate of which information is passed between them [177]. The classic example of this is the traditional Fitts' Law pointing experiment where a user is presented with several targets on a 2D display and input device *throughput* is determined by the function of their selection time and target difficulty [127, 245]. Amongst other things, good interaction is determined by how much time must be spent on a message or action [177]. For both AR and VR interaction, there are challenges in modeling human performance using Fitts' Law beyond target selection and standardised metrics for 3D manipulations are mostly missing [383].

Interaction as tool-use draws parallels from our adoption of physical or conceptual tools (such as a hammer or mathematical rules) to our engagement with computers [24, 31, 177]. Similarly, we can view AR holistically as a tool, however, it is more likely that if AR becomes more pervasive [142] we will instead view the virtual, physical, and augmented entities within AR environments as the actual tools at our disposal. For instance, in an AR manufacturing setting, a worker could use a screwdriver or a wrench to interact with virtual assembly instructions. Beaudouin-Lafon's 'instrumental interaction' [31] emphasises the importance of physicality and direct manipulation in virtual tools. Following this model, physical tools could not only be instruments of physical labour but also provide mediation between the User and virtual (*domain*) objects in AR environments. Section 2.5 expands on tools and interaction models further, but there are several key underpinnings of this interpretation of interaction that should be mentioned here. First, as described in Activity Theory, tools have the capacity to shape the user and how they act [218]. Additionally, there is a cyclical relationship between people, tools, and the tasks they engage in sometimes referred to as the 'task-artifact' cycle [66]. Generally speaking, tools influence people and vice versa. Interaction as tool-use also highlights the natural mediation between user and task that tools facilitate — "acting through the interface" [53]. Finally, this view of interaction is tied to *use* and to *tasks*, a confluence between the tool and the purpose it is used for — a tool and its job [24, 177]. To properly understand 'good' interaction, as defined by Tool-use, is to understand the canonical activities users currently and will endeavour to engage in using AR.

2.3.1 AR Tasks and 3D Interaction

Virtual environments, such as those created through AR, allow users to engage in a wide range of activities. These include direct **3D interaction** with virtual objects, **viewing or browsing** of overlaid information, and **system control** and symbolic input [43, 55, 221], all of which are facilitated by some form of interaction with virtual content. While this is not an exhaustive list of the tasks and potentialities of AR, this thesis primarily focusses on the physical interaction with virtual objects. However, it is important to note that AR interaction goes beyond physicalising virtual content, which we will further expand on in Section 2.4 and 2.5.2. The work presented in this thesis touches on aspects of **information browsing** and **system control** but predominantly explores **3D interaction** facilitated by physical tools [221].

One of the most appealing attributes of AR is the ability to render complex 3D models into physical space, allowing users to engage with virtual objects as if they were physically present in the room. As in the physical environment, without mechanisms to interact with objects in virtual environments, many tasks cannot be performed [221]. Therefore, AR interaction relies heavily on the field of 3D UI which provides interaction techniques to allow users to *select* and *manipulate* virtual objects — two fundamental tasks in 3D interaction and AR [55, 135, 221].

Selection tasks involve identifying/acquiring a specific virtual object from a set of objects available [221, 430]. In physical environments, this can be thought of as picking up or pointing at an object. In 3D interaction, there are a number of parameters to consider for selection such as the target distance [370], size, number, and position of distractors [394, 395], and the level of target occlusion [344, 427]. Selection in both VR and AR is a broad research topic in itself [118, 219, 242, 291, 307] and largely depends on the input modalities used, which we expand on in Section 2.3.2.

On the other hand, when 3D UI research discusses *manipulation* it is more precisely referring to spatial rigid body manipulation, i.e. manipulations that preserve the object's shape [216, 221]. As such, manipulation tasks are segmented into 3 sub-tasks: positioning (or translating), rotating, and scaling a virtual object [221]. Once again, there are parameters to consider for each sub-task mostly regarding the object's starting position/rotation/scale, the target position/rotation/scale, the required movements between them, and the required precision for each type of manipulation [221]. Importantly for AR, all 3 manipulations can each be performed in 3D which adds to the complexity when users are trying to precisely manipulate a model [39, 132, 287]. Precise manipulations, in general, can be difficult to achieve in Virtual Environments due to the lack of accommodation of combined sensory cues [112, 198, 227], something which we expand on in Section 2.4.

Beyond basic spatial rigid body manipulations, Virtual Environments can provide the means to manipulate the shapes and forms of objects, as can be done in desktop 3D modeling applications such as Blender [51] and Maya [259]. AR specifically has the capacity to extend these well-established desktop applications into physical space. There is a wealth of prior work on 'immersive' 3D modeling ranging from early work on immersive modeling [62, 99, 411] to more recent 3D sketching [18, 193, 322] and surface modeling tools [196, 261, 322]. One of the earliest Hololens demos showcased integration with Maya [267] and now even basic commercial modeling tools are available such as Google TiltBrush [382]. These demonstrate a clear vision for AR to play a fundamental role in future design and fabrication workspaces.

In immersive 3D modeling research, a distinction can be made between two main categories — 3D sketching (drawing lines and contours) and 3D modeling (creating polygonal or parametric surfaces) [322]. 3D UI research also distinguishes these types of manipulations from the essential actions of selection, positioning, rotating, and scaling [221]. Oftentimes, 3D modeling and sketching can be achieved by using the essentials of selection and manipulation combined with system controls to manipulate more abstract attributes such as object texture [221].

However, the more these modeling actions rely on system controls, such as menus and widgets, the less *direct* the interaction becomes. The same is the case if modeling interactions become heavily compartmentalised into the essential manipulations (selection, position, rotation, scale), resulting in more convoluted and laborious interactions. The required precision is also a key consideration when 3D modeling and sketching in AR [221]. We expand on current approaches to immersive AR modeling in Section 2.4. The AR tasks utilised in the studies presented in this thesis primarily focusses on 3D interactions, with Chapter 3 looking at 3D interaction more broadly, encompassing *selection*, *manipulation*, and *modelling*, whereas Chapter 5 focusses more specifically on **3D object manipulation**.

2.3.2 AR Input Modalities

Having reviewed the AR tasks and activities in the prior section, we will now examine the various input modalities and techniques for 3D interaction that have been developed for AR over time that are now commonly used. These input modalities can facilitate either **isomorphic**, a strict one-to-one mapping between the input and the output, or **nonisomorphic** interactions, a mapping beyond realism that would not be possible in a physical environment such as virtual hands or raycasting [132, 221]. Additionally, input modalities can also vary in directness, expressiveness, precision, and performance depending on the interaction technique.

Krueger's VIDEOPLACE is considered one of the first examples of early interactive AR systems [217, 274]. Though it functioned more like an interactive projection system, its use of techniques similar to those found in modern AR HMDs [266], such as mid-air deictic, manipulative, and semaphoric gestures [202] for interaction and content creation, heralding the development of more advanced AR technology. However, the system's limited 2D projection screen prevented it from providing the direct, 3D gestural interaction seen in modern Virtual Environments, which only became possible with advancements in display and tracking technology in the 1990s [22, 203, 409]. With the advent of computer vision, AR systems became capable of detecting user motion and gesture in real-time without the need for additional sensors [43]. This led to the development of a myriad of direct 3D gestures for spatial manipulation of virtual objects [60, 163, 234, 233]. In addition, wearables and biometric sensors can be used to detect user gestures when computer vision is not available or the hands are occluded, as demonstrated in the work of Lee et al. [231].

With a range of options for direct gestural interaction, Piumsomboon et al. defined a consensus set of gestures for virtual object manipulation using user guessability [308]. Subsequently, commercial AR headsets now employ a variety of standardized gestures such as pointing, grasping, pinching, and swiping to facilitate 3D interaction, either built directly into the AR HMDs or supported through additional integrated sensors such as UltraLeap [390]. Despite the prevalence of gestural interaction in AR, challenges remain, such as fatigue, precision, and expanding user capabilities beyond isomorphic manipulation. To address these issues, non-isomorphic gestures, e.g. virtual hands [54, 312] and raycasting [2, 292], or multimodal interaction [78] have been proposed and will be discussed further in Section 2.4.

Beyond mid-air gesture, device-based approaches have become popular over time, such as pointing devices, wands, and 6DoF (degrees-of-freedom) hand-controllers [30, 43]. The mouse is a fundamental and necessary tool for 2D applications on desktop computers, and it remains the most commonly used input device for 3D interaction and modeling. Previous evaluations have shown the mouse to work well for 3D manipulation [35, 429] and Krichenbauer et al. found no significant differences in performance between the mouse and direct 3D input devices [216]. Butterworth et al. used a handheld 6D mouse with a HMD, one of the earliest instances of a 3D controller, allowing for virtual object creation and editing when combined with a 3D UI [62]. Since then, hand controllers have become commonplace in virtual environments, especially for VR. Typically, hand-controllers are integrated with buttons and triggers, offering a lightweight solution for performing *system control*, similar to function keys in desktop systems [221]. For manipulation, hand-controllers can facilitate both non-isomorphic interactions, such as raycasting for selection [141, 165], and isomorphic interactions, such as prop-based or proxy-based manipulations where the controller acts as a direct substitute for the virtual object [168]. Controllers are a popular choice for indirect interaction with virtual interfaces and widgets due to their affordance for pointing and raycasting. However, related work has found that for 3D pointing and selection, controllers can have less throughput than the mouse [365, 377]. Finally, hand-controller form factor and sensing capabilities is a broad topic of research in itself and has strong overlaps with Tangible UI [178, 189] and Tangible AR research [46], which will be discussed further in Section 2.5.1 and 2.5.2.

Efforts to enhance AR interaction capabilities include combining different input modalities, such as hand gestures or controllers with natural interaction mechanisms such as gaze, speech, or bodily movement [43, 78, 279, 414]. The most researched approach is combining speech and gesture modalities, which have been shown to be intuitive for 2D and 3D interactions [88, 89]. Multimodal interaction can address challenges in 3D interaction prevalent in gesture and controller approaches, particularly for selection tasks. In AR/VR for example, there are instances where speech has been used for virtual object disambiguation during 3D selection and manipulation tasks [75, 404]. While research on multimodal input is well-established, it is not yet widely used in commercial AR systems and applications.

2.4 Interaction Challenges & Current Approaches

There is a prevailing notion that AR has the capacity to transform the way we work in physical environments [299, 337]. As virtual content becomes ubiquitous and embedded into our physical workspaces, users must be able to organise their virtual workspace as much as their physical. However, current AR systems face interaction challenges in facilitating **precise and sustained interaction** with virtual objects and, more specifically, 3D manipulations pose significant challenges for interaction designers [158]. There are three important aspects to enabling *precise and sustainable 3D interaction* with virtual objects in AR that have been explored in related work: 1) reducing the complexity of the task by compartmentalising different 3D manipulations or constraining degrees of freedom, 2) mitigating physical and cognitive fatigue as much as possible, and 3) providing feedback that accounts for our combined sensory cues. These approaches often overlap, as providing combined sensory feedback can also reduce task complexity and fatigue. This section will discuss these challenges and the related work that aims to address them, with a focus on enabling precise and sustainable interaction in AR.

As discussed in Section 1.2, our perception of physical environments is greatly enhanced by the combination and integration of various perceptual cues, such as spatial reasoning, object disambiguation, and perceiving affordance, which ultimately leads to a more accurate understanding of the environment. At a more granular level, when interacting with physical surfaces and objects with precision, our movements are guided by a strong sensory feedback loop, with mechanoreceptive feedback playing a crucial role in precise motor control [198]. Cutaneous mechanoreceptors provide high-quality neural images of object spatial structure and motion of the hand, enabling grip control and extending perception through tools and probes held in the hand through vibrations [198]. Commercial AR headsets have begun to incorporate some sensory cues such as visual disparity, shadows, shading, motion parallax, occlusion, and spatialised audio [266, 246]. However, both visual and haptic information is essential for precise interaction [112], and prior research has highlighted the over-reliance on visual perception in virtual environments [227].

In addition to sensory cues, another key challenge in AR interaction is reducing physical fatigue to engender sustainable interaction. As discussed in Section 2.3.2, the most common techniques for 3D interaction in AR involve hand gestures or controllers, which are typically used in mid-air. The '*gorilla-arm*' *effect* [188], first observed with the development of vertical touch screens and pervasive displays [150], occurs when users are required to extend their arms for prolonged periods of time without support, resulting in arm fatigue [195]. Similarly, mid-air interaction in AR can also cause fatigue with the '*consumed endurance*' model[164] showing that users can hold their arm extended at 90° for about one minute before experiencing fatigue, and likewise at 45° for about two minutes. To address this issue, several techniques and guidelines have been proposed, primarily focussing on designing interactions that support or account for the natural resting positions of users.

2.4.1 Reducing complexity in 3D Interaction

Reducing the complexity of 3D interaction typically involves **dividing** the labour, **constraining** the environment or interaction, or **expanding** user function. As mentioned in Section 2.3, direct manipulation in Virtual Environments affords ease of learning and use and the ability to perform simultaneous complex 3D manipulations in a single operation as we would in physical environments [155]. The benefits of direct manipulation can be maintained even as the required precision of the task increases by leveraging the benefits of *bimanual manipulation*.

Classic work from Buxton and Myers on bimanual manipulation found that users were able to perform tasks in parallel and even split tasks between two hands, demonstrating a significant performance increase over one-handed techniques [63]. This has been applied in AR/VR with Song et al. proposing a *symmetric bimanual* handlebar metaphor to allow 7-DoF manipulation of virtual objects [353]. Users performed rapid translation, rotation and 1D scaling of a virtual object by manipulating a handlebar that intersects an object. There are also examples of *asymmetric bimanual* techniques based on Guiard's work on hand asymmetry [144]. Stoakley et al. proposed *World in miniature*, one of the first instances to utilize an asymmetric bimanual technique in a virtual environment, with the dominant as the primary means of navigation and the non-dominant used to reference a miniature world [358]. Hinckley et al. further expanded on this with labour division across different input modalities controlled by the dominant and non-dominant hands with 'Pen+Touch' — pen writes and touch manipulates [170]. While these techniques demonstrate the benefits of bimanual direct manipulation, precision can still be improved when performing 3D interactions in AR. Chapter 5 of the thesis tackles precise interaction techniques directly and explores the role of bimanual techniques in affording precision in 3D environments.

Moreover, labour can also be divided by leveraging other input modalities to support manipulations in Virtual Environments. Gaze, in particular, is a popular modality to combine with mid-air interaction, for example, Yu et al. combined gaze and hand features to aid in 3D object manipulation which proved to be useful when interacting over large distances [425]. Pfeuffer et al. achieved something similar with '*Gaze+Pinch*' [306] and Mardanbegi et al. used gaze alignment to a virtual object through a virtual tool to edit and manipulate the objects properties [255]; both follow the principle of "gaze selects, hands manipulate". In this way, direct mid-air interaction is only called upon when necessary while leveraging our effortless and natural gaze behaviour to disambiguate targets.

Another approach to increasing precision is by providing environmental constraints to interaction, for example constraining the type of manipulation performed by direct interaction or reducing the degrees of freedom (DoF) of a manipulation [264, 265]. 'Snap-dragging' is a technique inherited from 2D interfaces, which Brier generalised into 3D [40]. Others have explored pre-defined constraints for 3D manipulation which reduces the DoF available for a user to manipulate [288, 350, 351, 366]. However, this work is 3D applications for desktop environments, falling back on menus and widgets to control object translation, rotation, and scale or reduce the manipulable axes, which is not immediately transferable or desirable for Virtual Environments. Hayatpur et al. attempted to preserve direct manipulation in VR by allowing users to implement constraints by micro gestures on a hand-controller [155]. The proposed techniques utilized shape constraints to separate DoF for precise object alignment and manipulation. However, there were several limitations primarily related to fatigue and usability. These were in part caused by performing the required gestures whilst simultaneously balancing the controller in the same hand [155]. Additionally, some parts of the technique also required users to keep their hands suspended in mid-air leading to fatigue issues after prolonged interaction.

Altering the C/D (control-display) gain has been a popular technique for expanding user function and increasing precision on desktop and touch display screens [7, 50, 68, 375]. The early work of Poupyrev et al. found that for 3D rotational tasks amplifying the rotation decreased task completion time without affecting user accuracy [316]. Other work has also looked at adapting C/D gain for Virtual Environments such as *PRISM*, a dynamic C/D gain interface for mid-air gestural interaction in VR [131]. The C/D gain was implicitly increased and decreased based on the user's hand movement speed when interacting with a virtual object [131]. The design of *PRISM* was based on the principals of Fitts' Law and utilized the user's natural behaviour to adjust the C/D gain, akin to mouse/cursor control [127, 245]. Furthermore, Poupyrev et al.'s '*Go-Go*' interaction technique [312] attempted to improve on basic virtual hands by allowing users to increase the length of the virtual hand exponentially (i.e. the C/D gain) by stretching out their physical hand beyond a distance threshold. Since then there have been a number of virtual hand techniques that build on this, but all primarily focus on selection [29, 313, 334]. Implicit C/D gain control has also been incorporated into techniques that provide manipulation constraints such as Wigdor et al. with their '*Rock and Rails*' technique for touch screens [413], which was later expanded on by Hayatpur et al. for 3D manipulations [155]. Despite accounting for user hand instability and reach, applying dynamic C/D gain to mid-air techniques means that while users can more readily control their precision, they can still encounter the same issues of '*gorilla-arm*' and fatigue when performing prolonged interactions. Chapter 5 explores this further and provides several example interaction techniques demonstrating how dynamic C/D gain can be leveraged to enhance user precision in 3D AR interactions.

Like with virtual hands, user function can be expanded by manipulating other aspects of the virtual representation of their body in the virtual environment.'*Beyond Being Real*' is a framework that leverages the plasticity of the human sensorimotor system by manipulating the virtual representation of the user's body in VR to perform movement-based interactions that would otherwise be impossible in the physical environment [5]. For example, multiplying the number of virtual hands [334] or adapting the proportions and limbs of virtual avatars [262]. However, a big constraining factor of this framework is that it relies on having complete control over the user's perception of their own body to remap motor tasks. VR provides this inherently, however, for AR this proves more difficult as users are simultaneously grounded in the physical environment through their bodies. As a result, these *beyond being real* interactions [5] are more difficult to achieve in AR due to challenges such as physical-virtual body misalignment [81].

2.4.2 Accounting for combined sensory cues in AR

A number of approaches have been proposed to provide crucial haptic cues when performing precise 3D interactions in Virtual Environments ranging from wearable and situated haptics, on-body interaction, leveraging physical surfaces, and tangible interfaces and artifacts.

Firstly, wearable devices are a common aspect of VR and AR research dating back as early as the late 70s with wired gloves used for sensing interaction [52, 363]. Nowadays wearable devices are not only used for sensing but also for simulating haptics of virtual objects [94]. Haptic simulation in Virtual Environments typically falls into two categories *passive* haptics, physical objects which provide feedback to the user simply by their shape, texture, or other inherent properties; and *active* haptics, feedback that is actuated and controlled by a computer [94, 238]. Wearable haptics in VR and AR typically utilise a form of active haptics to simulate the properties of virtual objects or surfaces such as texture, temperature, size, form, stiffness, and weight. Several works have looked at vibrotactile feedback in wearables to simulate texture [205, 309, 417], while others have explored pneumatics [378] and hydraulics [149] for texture and temperature.

Additionally, pneumatic wearables have been explored to simulate virtual object form and stiffness [378], with other work exploring exoskeletons and electrostatic brakes to simulate stiffness and form [82, 83, 115, 166, 237]. Skin deformation has also been used in wearable devices for simulating the weight and inertia of virtual objects [82]. A considerable amount of wearable haptic research focusses on simulating haptics in the hands but there are many examples of simulating haptic cues elsewhere on the body such as the head [205], arms [69], and full-body [214]. Beyond research, there is also a growing number of commercially available haptic gloves [305] such as HaptX [151], CyberGrasp [95], and Dexmo [101].

Wearable active haptic devices can provide crucial mechanoreceptive feedback required for interaction [198], but there are a number of issues when considering their application in AR. Firstly, most haptic wearables focus on simulating virtual object properties, such as texture, form, and stiffness, assuming users will be directly interacting with virtual objects ultimately depending on mid-air gestures. As a consequence mid-air fatigue persists and, in some cases due to the added weight of the wearable device, can be exacerbated. Additionally, for AR specifically, wearable devices especially on the hands act as a barrier between the user and the physical environment. If the user is to engage in physical work while in a Virtual Environment, a wearable haptic device could hinder their movements and engagement with the physical environment. Wearable devices also only provide feedback for the user that is wearing them, so the scalability of these devices to multiple users is also something to consider for AR.

An alternative to wearable active haptics is *situated* active haptics in the physical environment. The main difference between situated haptics and wearables is that the device simulating the feedback is '*encountered*' by the user in the environment and generally the user's hands are completely unadorned [422]. Methods for generating situated haptics range from using ultrasound, acoustics, and even using drones to dynamically rearrange into props and surfaces. UltraHaptics is one of the most notable examples of providing localised, mid-air haptic feedback above an interactive surface using ultrasound [67], something which has been expanded on by others [241, 258]. Sodhi et al. achieved a similar thing but instead used pneumatics to generate '*free air textures*' as a haptic display [352]. Other work has utilised acoustics to programmatically levitate passive props above a surface and generate discrete points of haptic feedback [171, 263].

In a similar capacity instead of levitating props other work has opted to use drones to dynamically construct objects and surfaces for the user to encounter in the virtual environment [3, 56, 137, 175, 422], but this is primarily implemented for VR where the drones are not visible. Generally, a constraint of these types of situated haptics is that, while dynamic in terms of where they can be generated, the force feedback they provide is relatively low. Alternatively, other research has considered shape-changing displays and dynamically moving props to provide a stronger sense of force feedback through more rigid materials [4, 13, 80], but these displays are much more limited in terms of dynamic placement and movement.

Beyond active haptics, researchers have considered leveraging the passive haptics that is available on our hands and bodies as an input and output platform [37, 153]. There have been numerous techniques developed for menu and widget-based interaction in AR/VR utilising user's arms, hands, and palms as projection spaces for text entry and object selection [20, 120, 197, 426]. Furthermore, Pei et al. introduced the concept of hand interfaces in which the users mimic the form of a virtual object to have their hands *become* the object in VR which they then can manipulate using their hand as a proxy [304]. In general, the exploration of on-body interaction is limited to widget-based interaction or used in combination with other modalities and there are few instances of using the passive haptics of the body for direct 3D manipulation aside from the work of Pei et al. [304].

Another example of leveraging passive haptics is research around augmented workbenches and tabletops combined with HMD AR, which has a substantial history serving as a foundation for many AR applications [34, 199, 204]. Flat surfaces in particular are important for enabling collaborative work providing a shared interaction space [336], allowing the placement of physical props [161, 302], and supporting precise interactions that are much more difficult to achieve in mid-air [322]. This body of work also heavily overlaps the extensive research conducted around interactive table tops specifically around object manipulation on-surface and above-surface.

Interactions in the air was a technique to allow users to seamlessly switch between interacting with a tabletop to manipulating digital content in 3D above the surface [162]. *Mockup Builder* made advancements on this concept, following Guiard's asymmetric bimanual model [144] to allow for interaction continuity when editing 3D models with one hand operating the tabletop and the other operating the virtual object in mid-air [98]. *SymbiosisSketch* combined free-hand and surface-based sketching to create virtual objects combining tablets, pen interaction, and HMD AR [18]. Additionally, similar techniques and applications have been designed for fabrication, such as *MixFab* [406] and *DualCAD* [271] combining desktop interaction, mobile devices, and HMD AR. The projects discussed so far have highly decoupled input and output spaces between the interactive surfaces/devices and the AR projection/interaction space. The work of Reipschläger et al. attempted to blend virtual environments and interactive tabletops, even more, using the virtual environment as a direct extension of 3D modelling software running on the tabletop [322]. Combining AR with surfaces is just one example of how the physical environment can be complimentary to 3D interaction and even provide *opportunistic* *controls* [157] something which we expand on in Section 2.5.2.

Finally, another approach to providing haptic sensory cues and a primary focus of the thesis is Tangible User Interfaces (TUI) and objects [338]. Tangibles have the ability to utilise both passive and active haptics and have widely been incorporated in virtual environments to produce rich haptic experiences [46]. There are several examples of active haptics being incorporated into existing VR and AR hand controllers and new form factors to facilitate new haptic experiences [33, 84, 240, 340, 367, 412]. However, these tangible devices are primarily concerned with simulating haptic sensations of virtual objects in VR, whereas Tangibles can also be used to enable new forms of interaction beyond direct gestures. Additionally, Tangible objects can have spatial relevance in both the virtual environment and the physical environment, acting as somewhat of a mediator or 'handle' between the two. Section 2.5.1 and 2.5.2 will further expand on Tangible User Interfaces and Tangible AR.

2.5 Next-Generation Tools and Interfaces

Now we have discussed 'AR' and 'Interaction' in the context of the thesis, let's now discuss '*Tools*'. Tools and Tool-Using are complex concepts to define and are subject to considerable debate [343]. Activity Theory posits, amongst other things, that physical and conceptual tools, such as instruments, procedures, and methods of labour, are '*integral and inseparable components of human function*'. Additionally, it suggests that human behaviour is not directed 'from the inside', based on biological urges, but 'from the outside', based on context and the use and creation of artifacts [218]. For the purpose of this thesis, we follow a definition that is not necessarily encompassing of all types of tools and tool-use, but one that is more focussed on the subject matter in the thesis: '*A naturally occurring or designed physical object*'— **Tools** —'*that is utilised for an immediate purpose*'—**Tool-Use** [275]. As mentioned in Section 1.1, appropriating this definition for Tools in AR requires some understanding of what the '*naturally occurring or designed objects*' are and what the '*purposes*' or functions are.

Starting with '*objects*' in AR, these could be every day (naturally occurring) objects in our physical space that are appropriated for interaction [157], or bespoke objects and controllers designed for a specific '*purpose*' or interaction. Concerning '*purpose*', as mentioned previously in Section 2.3 we have some understanding of what the activities might be in AR — 3D Interaction, browsing overlaid information, system control. These tools can enable mimicking of the real world, for example providing tactile and haptic cues for virtual objects [94], but also they can expand user function in AR environments facilitating interactions that are not possible in physical space — going beyond reality [5].

As mentioned in Section 1.1, HCI research has a rich history of developing tools for mediating with technology leading to a number of taxonomies and frameworks providing a foundation for exploring AR tools. Bødker's seminal work describes how users operate '*through*' an interface on objects of interest (such as documents) and, with good design, the interface becomes hardly noticeable [53]. Beaudouin-Lafon expands on this for 2D desktop interfaces describing interaction on (domain) objects using 'instruments' [31], virtual tools that yield new styles of interaction based on the principles of direct manipulation [185]. *Instrumental Interaction* was then extended by Coutrix and Nigay for 'Mixed Reality' [91] (virtual environments), considering 'mixed objects', physical and digital tools or objects of interest, and other interaction modalities such as Tangible interaction. More recently, Jacob et al. proposed '*Reality-based*' interaction a framework that unified interfaces that more closely correspond to daily practices in the non-digital world [194]. For instance leveraging people's knowledge of the physical world, bodily and environmental awareness, and social awareness for interaction.

Ultimately, interfaces that meet humans where they excel or are physically limited will have the capacity to expand function. Considering this framework, the field of Tangible User Interfaces (TUIs) is the closest example of physical tools being utilised for digital interfaces and has widely been explored for AR.

2.5.1 Tangible User Interfaces

Tangible User Interfaces (TUI) are a post-WIMP interface type that emerged in the 1990s and the surrounding research is concerned with providing physical representations of digital information and controls [338]. The preliminary work of Ishii, such as *Bricks* [128] and *TangibleBits* [189, 191], underpins a large portion of current research and application domains [45, 129, 199]. The concept of graspable UI explored the manipulation of digital content using physical clutches, utilising the advantages of bimanual interaction and spatial reasoning for collaborative workspaces [12, 144, 169, 210].

There are a number of application areas for TUIs such as learning and edutainment, problem-solving and planning, and information visualisation [338]. TUIs have also been leveraged in rapid interface prototyping, for example, Kelly et al. [206] and Zheng et al.'s [431] work produced popular TUI elements, such as knobs and sliders, using low fidelity materials and fiducial-based computer vision tracking. Beyond low-fidelity prototyping, there are tabletop approaches that also use external computer vision tracking to build toolkits and development platforms for TUIs such as reacTable [199] and reacTIVision [201], Madgets [407], and SLAP Widgets [408].

Previous work has provided taxonomies for describing and analysing TUIs utilising concepts such as embodiment, metaphor, and representation [125, 126, 338, 387]. Fishkin provided an initial two axes taxonomy for TUIs using levels of embodiment and metaphors based on existing systems at the time [103, 125, 387], however, our understanding of embodiment has significantly evolved [212, 355, 392]. Metaphors enable TUIs to become more "analogous to the real-world" denoting what an interaction device *is* — shape and form — and what *can be done* with it — motion and input [125]. To this end, a physical object's power of representation can be utilised in tandem with a form factor to facilitate more complex tasks in AR, such as 3D manipulation and modeling, and further diminish "cognitive seams" in interaction [190, 251].

Another interesting dimension to TUI is the concept of object affordance — how an object's "*physical form and manipulability convey how to handle it*" — coupled with prior knowledge of the "*everyday physical world*" [136, 178, 281, 283]. Norman extended the definition to include 'perceived' affordances, typically visual cues in the interface relying on cultural conventions, and 'real affordances' that exist independently of user perception [283]. Leveraging the affordance of physical objects is a common sentiment in TUI literature, utilising a user's "*haptic interaction skill*", "*expectations of the real world*", and "*familiarity of everyday physical objects*" [44, 159, 178, 189, 194]. However, there have been many criticisms of the concept advocating that affordances and tangible interaction, in general, rely on "*spontaneous reactions*", circumventing "*conscious decision-making*" without means of "*recovery and reflection*" [178, 318]. Further to this, Hornecker explains how object affordance, properties, and functionality are not universally interpretable even going unnoticed [178].

This thesis utilises a functional characterisation of object affordance, which posits that *the physical form and manipulability of a tangible object convey the appropriate way to operate it for a given task and context of use* [136, 178, 281, 283], which we elaborate on further in Section 3.1. As mentioned prior, object affordance is subject to individual, societal, and cultural influences, and therefore highly subjective and dependent on a user's or designer's interpretation. While the **mapping** of affordances to digital interaction is desirable it is "*far from straightforward*" [119, 178], and combined with the challenges of 3D interaction [158], it can be difficult to get the mapping 'right' for physical tools in AR. To address these issues, we combine affordance mapping with guessability methodology and user-elicitation [419] as an exploratory approach to AR Tool-Making, that aims to develop novel interaction devices and techniques through the appropriation of physical objects.

On the subject of gestures, tangible objects can be combined with gestural interactions presenting novel design opportunities for 3D interaction. According to Van den Hoven et al., "*gesture theories and classification from social sciences have not explicitly addressed gestures with objects in hand*" [180]. Previous work has explored so-called *tangible gesture interaction* [10, 180] with each focusing on different types of objects such as mobile devices [167, 324], controllers and props [229, 398], custom tangibles [90, 434], and environmental objects [86, 159, 421]. In light of the wealth of work on tangible gestures, there is an opportunity to explore similar techniques for typical AR activities such as 3D transformation, modeling, and editing.

2.5.2 Tangible AR

Succeeding the body of work on TUI, Tangible AR aimed to combine the enhanced display properties of AR HMDs with intuitive manipulation afforded by physical objects. Billinghurst et al. describe how TUIs at the time were limited in supporting 3D viewing and interaction with virtual objects, often with a disconnect between task and display space [46]. Billinghurst et al. proposed an independent classification of Tangible AR interfaces and later defined a taxonomy across two domains: Space-Multiplexed and Time-Multiplexed interfaces [46]. In Space-Multiplexed interfaces each function has a physical device occupying its own space which was exemplified by the initial prototypes of *Tiles* [315], *Shared Space* [47], and *ARgroove* [46]. In Time-Multiplexed interfaces on the other hand a single device controls different functions at different points in time as showcased in the *VOMAR* system [46]. These interfaces were an early example of how new metaphors and forms of interaction were enabled when AR tasks were mediated by physical tools, utilising passive haptics to reduce the 'functional' and 'cognitive' seams between interaction and display space [45, 190].

More recently, the physical form factor of tools and controllers in AR has been explored. Several works explore and evaluate basic primitive shapes, such as cubes and spheres, and even compare these form factors against state-of-the-art AR controllers [111, 192, 228]. Englmeier et al. showed how a static spherical AR proxy had significant advantages over asymmetric controllers for 3D manipulation regarding task-completion time and workload [111]. In addition, the benefits of cubes as tangible interaction devices have previously been explored [235, 339]. Particularly the work of Lefeuvre et al. [235] categorises the distinct affordances and properties of cubes by "manipulation, placement in space, arrangement, multi-functionality, randomness, togetherness, physical qualities, containers, and pedestal for output". Issartel et al. explored these cube properties applied as a tangible window into a virtual environment allowing for a range of 3D interactions [192], and Lee et al. looked at a bimanual cubic interface with two different interaction techniques inspired by object assembly [228]. There has also been considerable work designing more bespoke and familiar physical tools for 3D interaction such as chopsticks [423], paint brushes [364], and hammers [359]. Arisandi et al. presented a toolbox of physical tools for 'virtual hand crafting', such as a hammer, tweezers, and knives, a more recent example of a Space-Multiplexed interface and objects that more overtly represent traditional, non-digital tools [16].

Other research has looked beyond using physical objects in AR purely for their passive haptics but instead looked at the creation of interactive artifacts. For example, Savage et al. focused on interaction with physical objects while utilizing a computer-vision-based tracking approach to detect user gestures and create 'active' objects [331]. These artifact-based interfaces are commonly used in virtual environments due to their manipulability and ability to be used as a proxy for virtual objects [17, 111, 117, 273, 434]. The work of Feick et al. [117] produced a toolkit for prototyping such tangible artefacts for virtual environments, using low-fidelity and modular 'shape primitives', to support proxy-based interactions. Beyond proxy interaction, the work of Angelini [10] and Van den Hoven et al. [181] shows how combining a tangible artifact with hand gestures, both motion and surface-based, can effectively recreate many common UI elements (knobs, sliders, buttons) using a single artifact.

In addition to bespoke designed tangibles, previous work has shown how AR can be used to explore ad-hoc and nomadic interfaces through everyday objects. Walsh et al. developed an architecture to allow tangible interfaces to be defined ad-hoc through an abstracted set of interactions [403]. Furthermore, 'everyday' physical objects can be seen as an opportunity to facilitate Tangible AR interactions [157, 159]. *Annexing Reality* was a system developed to opportunistically annex physical props from a user's environment to create rich haptic experiences when interacting with digital content, attempting to blend the virtual and the physical [159].*Affordance++* utilised active haptics, specifically force feedback via electrical muscle stimulation, to allow instrumented 'everyday' objects to communicate their affordances through interaction or even augment the affordances for virtual content manipulation [243, 244]. Further to this, Gupta et al. also found the instrumentation of physical objects to be a "*strong direction for AR applications*" [146].

2.5.3 Building Modifiable AR Tools

Billinghurst et al. highlight the difficulty of dynamically changing the physical properties of TUIs [46], and Jacob et al. discuss the trade-offs between the versatility of GUIs and the rigidity of TUIs [194]. However, recent research has investigated the role of physically-modifiable objects as proxies for virtual models in AR and VR [17, 117, 273, 434]. For example, Feick et al. presented a rapid prototyping toolkit, made of shape primitives, for creating object proxies with moving parts which were preferred by users over traditional controllers [117]. The work of Arora et al. was similar but instead used Lego bricks over shape primitives [17]. Arif et al. and East et al. demonstrated how smartwatches could be re-purposed as modifiable tools for interactive tabletops [15, 107]. Spatially reconfigurable tangibles have also been explored for tabletop interaction, for example, Nowacka et al. introduced autonomously moving objects with sensing capabilities for interaction on tabletops [286]. Le Goc et al. similarly and more recently presented a tabletop swarm interface of tangible microbots that could self-reconfigure based on users' input and application scenarios [223, 224]. Tangible tools for AR such as these have the capacity to become more 'versatile' if they are designed to be not only virtually modifiable but also physically modifiable.

Additionally, the field of shape-changing interfaces contributes to the development of adaptable and dynamic form factors [87, 133, 319, 362, 320], which can inform the exploration of physically-modifiable tools in AR. The work of InFORM [130] and EMERGE [372] are perhaps archetypal examples of shape-changing displays, predominantly in-situ and supporting direct and indirect user manipulation. Variations of the situated shape-changing display have been developed such as CairnFORM [96], Relief [236], and more recently STRAIDE [110], but shape-change as a concept is much broader. To give an example, the work of Roudaut et al. showed how mobile devices can adapt their shapes dependent on the context of use to offer improved affordances [327]. There are further examples of modifiable interfaces that leverage elements of shape-change to provide users with flexible and reconfigurable interaction devices [71, 148, 182]. More recent work has looked at combining shape-changing interfaces with Virtual Environments, for example enhancing the appearance and visual effects of shape-changing objects using Spatial AR [239], providing modular shape primitives to build bespoke proxy objects for VR [117], or to generate new types of mediated haptic experiences of virtual objects [220, 260].

In addition to exemplar prototypes, the field of shape-change offers several notable contributions in terms of classification. One such contribution is the work of Rasmussen et al. [320], which provides valuable insights applicable to physically-modifiable tools. Their classification encompasses different types of change, such as form, volume, and texture, and introduces a transformation vocabulary that describes the process of material change using factors like velocity, path, or space. Additionally, Taher et al. [373] offers a technical approach to classifying shape-change, focusing on the methods of computer-based actuation, such as electro-mechanical or magnetic mechanisms. They consider factors like actuation speed, granularity, force, size, and complexity of control. Moreover, Sturdee and Alexander [361] propose a different classification for shape-change prototypes. Particularly relevant to the work presented in Chapter 3 is their characterisation of the **physicality** of the interfaces, as they found it to have the most influence on user interaction. These diverse classifications, along with the associated challenges in shape-change highlighted by Alexander [8], provide a basis for the framework presented in Chapter 3 for ideating and creating physically-modifiable AR tools. More specifically, we are considering the **physicality** [361] of existing objects and how these attributes can be leveraged for interaction with virtual objects in AR.

2.5.4 Tracking, Sensing, and Prototyping

Considering the related work discussed so far, we can describe two general approaches to tracking tangible interaction: (i) using an external tracking system or (ii) instrumenting an object or user. Both approaches have various limiting factors. For example, external tracking systems, such as computer vision and infrared lighthouse tracking, can be disrupted by occlusion when tracking users, physical objects, or gestures [19, 173, 424]. Additionally, instrumenting a user can hinder their engagement with the physical environment and absolute tracking of instrumented objects can be difficult. There are also examples of combining external and instrumented tracking methods to mitigate both of their limitations [183].

Considering sensing interaction on instrumented objects, capacitive sensing is a common and well-documented approach, with particular benefits for rabid fabrication and prototyping [61, 140, 289, 332, 335]. The work of Schmitz et al. [335] and Burstyn et al. [61] explored capacitive input on 3D printed interactive objects, specifically leveraging conductive filament and dual-extrusion printing, to create discrete touch-sensitive areas on any 3D object. Capricate [335] in particular provided tools for designers to modify virtual models to be printed and instrumented with touch-sensitive areas. Beyond 3D printing of capacitive objects, so-called 'loading mode' capacitive sensing has been adopted for gesture recognition [140]. Commonly, a number of capacitive areas are used to form capacitive sensor arrays to track touch across a designated area over time [139, 276, 346, 357]. Nelson at al. [276] incorporated 4 to 12 interactive areas to detect gestures on fabric and provided an initial exploration of different capacitive plate shapes and designs.

For prototyping tangible interaction, toolkits have been proposed as a productive approach to mitigate engineering challenges in building interfaces [226]. These toolkits range in levels of fidelity from paper prototyping [206, 431] to sophisticated electronic toolkits [17, 117]. A common evaluation approach for understanding the feasibility and generalisability of toolkits is to actually use the toolkit to design and develop a number of demonstrative applications [179, 226, 256].

2.6 Summary

The related work on Augmented Reality suggests that it is becoming increasingly prevalent and integrated into daily work practices. A key aspect of AR is 3D interaction, such as manipulation and modeling, which is crucial for interacting in virtual environments and conducting work in AR. However, there are challenges with 3D interaction in AR, particularly in terms of prolonged interaction and precision. Research is being conducted to address these challenges, including the use of physical tools and tangibles for AR, which form the foundation for this thesis. So let us summarise the knowledge we can gain from related work for *Tool-making*, *Tool-modifying*, and *Tool-Using* in AR.

Tool-Making: As Napier describes — "*Imagination is basic to tool-making*" [275] — and we must first conceive of potential tools for AR interaction. Combining AR and TUI has shown to be promising in previous work, simultaneously leveraging the intuitiveness, materiality, and spatiality of tangible interaction, and the enhanced display capabilities and 3D interaction afforded by AR. To explore tool-making in a systematic manner, we can draw on the existing classifications in shape-change [320, 361, 373] and wider methods in HCI such as guessability methodology [418, 419], affordance mapping [178], and "tangible gesture" [10] from previous research. However, most work on tangible interaction has focused on passive objects, leaving the area of shape-changing tools relatively unexplored and lacking methods for description, classification, and comparison.

Tool-Modifying: Considering the related work, our objective is to enhance and expand upon previous tangible prototyping toolkits [226] with an exploration of physically modifiable tools for AR interaction. We aim to build upon previous instances of tools developed for AR, including those employing basic shapes [111] or custom form factors [117]. Our approach is to merge research on fabrication techniques of 3D printed interactive objects [335] with instrumented artifacts and external tracking techniques [183], to create a versatile tool that can be applied to a variety of AR applications.

Tool-Using: After reviewing the current state-of-the-art interaction methods for AR systems [14, 264, 280], we have identified limitations with the widespread use of hand gestures and hand controllers. Based on previous research, we have determined that tangible objects can offer unique opportunities for interaction, specifically by leveraging object affordances to enable new interaction techniques and applications. These new interactive artifacts present unexplored design spaces that require further examination to understand their full potential. Additionally, by comparing these new tools to current state-of-the-art interaction methods, we can gain a deeper understanding of their benefits and limitations.

Based on the previously examined related work, this thesis presents a fresh outlook on the future of interaction in Virtual Environments, with a specific focus on the making, modifying, and using of physical AR tools. Just as AR has the potential to bridge the gap between reality and virtuality, the tools used to facilitate interaction and tasks with virtual objects should also diminish the divide between the user and the tool designer. This approach aims to provide users with increased flexibility and autonomy in conducting their work within these blended environments while allowing for deep personalisation of their tools, similar to the level of customization available to crafters and users of physical tools for the physical world.

In the next chapter, we delve into an approach for tool-making that leverages various methods for interaction design, such as user elicitation. In this way, we aim to actively engage users in the process of ideating tools and provide them with tangible references in the form of preexisting physical objects. These objects serve a dual purpose: as handles to facilitate the design process of AR tools and as a means to explore the mapping of affordances. Through the empirical studies described in the following chapter, we aim to identify any potential generic design guidelines that arise through affordance mapping to inform the physical and virtual characteristics of future AR tools.

Chapter 3

Tool-Making

This thesis aims to explore physical tools for Augmented Reality (AR), specifically 3D interaction with virtual objects. We begin by tackling the first research question around *Tool-Making* in AR (**RQ1**): *How do we ideate tools for Augmented Reality that are both physically and virtually modifiable?*. In HCI, the creation of tools is a long-standing practice of producing a variety of interaction devices for different technologies using established methodologies [222]. Additionally, tangibles have been employed to create more immersive and experiential interactions that are anchored in the physical space, leveraging metaphor and object affordance in different contexts[191, 338]. In our work, we focus on physically-modifiable objects (PM-Objects), i.e. objects that have distinct and innate properties for shape change and reconfiguration based on user actuation, with the goal to create more '*versatile*' tools [194] that enable more expressive and precise interaction — physically-modifiable tools (PM-Tools). To answer **RQ1** we pose three additional sub-questions:

- **RQ1.1:** How can we categorise and compare PM-Objects based on their relative properties and affordances?
- **RQ1.2:** How do PM-Objects with different properties and affordances map to complex AR tasks such as 3D manipulation and modelling?
- **RQ1.3:** What are the commonalities in the way users apply metaphors and couple input and output when using a PM-Object for a specific AR task?

In this chapter, we start by motivating some of the materials and methods we employ in our exploration of tool-making. We then describe the design space that emerges when specifically considering PM-Objects and present the five core steps to our methodology. These are based on various pre-existing methods such as user-elicitation and guessability [419], affordance mapping [178], and 'tangible gesture' [10, 180]. The chapter continues by introducing the first two studies of the thesis, a group design workshop and user-elicitation study, and presents the findings and the accompanying discussion. Finally, the chapter concludes by synthesising the findings from the two studies into conceptual examples of physically-modifiable tools (PM-Tools) and reflects on the presented methodology and framework.

3.1 Motivation

In our work, we explore tangible interaction devices to determine if properties and affordances can be leveraged differently for complex tasks in AR such as transformation and modelling virtual objects. In light of this, there is a combination of advantages for physical objects and tools highlighted in TUI literature and state-of-the-art AR interactions [191, 338, 387]:

Passive Haptics: As with all types of tangible interaction and much like hand controllers, tangible objects have innate tangibility in the form of passive haptics. For interactions that require a level of precision, this provides tactile motor-control feedback which is necessary for precise actions and beneficial for proprioception [21, 415].

Spatial Permanence: Tangible objects are individual artefacts that can be placed or positioned in different parts of a user's environment. For AR environments, an object's spatial permanence can be leveraged for distributed, state-based, and localised interactions without the need to carry the object with you at all times [12, 45, 194, 210].

Affordance: The physical form and manipulability of a tangible object convey how the object should be operated [136, 281, 283] and is informed by our preconceived notions of the object, the physical world, and the cultural context [178]. In the context of AR interaction, physical affordances can be used to scaffold complex virtual interactions and demonstrate the manipulation constraints of a virtual tool or object. This can be particularly effective in supplementing more abstract interactions, such as manipulating colour or deforming the mesh of virtual objects, with tangible objects providing cognitive mediation [178, 191, 281, 283]. By leveraging the physical affordances of tangible objects, users can gain a better understanding of the virtual affordances in a virtual environment.

Discrete Input: A subset of tangible objects have a distinct area or mechanism for input that can span multiple input modalities, i.e. a button press or a dial rotation [71, 148, 182]. Similar to 3D hand controllers, these input methods can be leveraged for **system control** enabling precise refinement, variable manipulation, or clutch-based interactions [125, 338].

For this work we focus on Head-Mounted Display AR. When considering how to explore PM-Objects and AR, there are a plethora of Mixed Reality mediums that could be adopted such as handheld AR, projection AR, fish-tank Virtual Reality, and HMD VR. However, there are a number of requirements to allow users to explore PM-Objects unhindered. First, the user's hands should be completely uninhibited when interacting with physical objects allowing for bimanual interaction if required. Second, they should be able to explore both forms of direct and indirect interaction between the physical objects and the virtual models. Third, they should be able to leverage the physical environment unhindered, such as surfaces, if they see fit for a given interaction. These requirements rule out most mediums aside from projection and HMD AR, however, projection AR suffers from shadow and occlusion issues when directly interacting with a virtual model. While we distinctly focus on HMD AR, we aim for exploration of PM-Objects to have transferrable insights for other mediums of AR. We also focus on a subset of fundamental AR tasks and activities — 3D manipulation and modelling of virtual objects.

3.1.1 Using Non-Digital Physically-Modifiable Objects

In addition to exploring the general benefits of using tangibles in AR, we focus on a subset of non-digital tangibles called physically-modifiable objects (PM-Objects). These objects have the same advantages as basic TUI, while also having the ability to be deformed or reconfigured allowing for new forms of interaction in AR.

Controlled Deformation: A subset of physical objects that have inherent material properties meaning that they can be compressed, expanded, bent, and twisted into a variety of shapes. For example, clay can be compressed and moulded into any shape, or an elastic band can be stretched and relaxed to its original shape.

Configurability: A subset of physical objects that can be changed through the construction or deconstruction of discrete components such as Lego or magnetic cubes, or reconfigured via a shape-changing mechanism, such as a Hoberman sphere or Rubik's cube. Configurability differs from deformation as it is not associated with the materiality of the physical object.

We investigate the potential of using non-digital, pre-existing PM-Objects such as toys, gadgets, and fidgets that can be deformed, reconfigured, and have mechanisms for input in AR. We choose to use these non-digital objects for two reasons:

1. They allow for exploring user-designed gestures without artificial constraints introduced by instrumented objects.

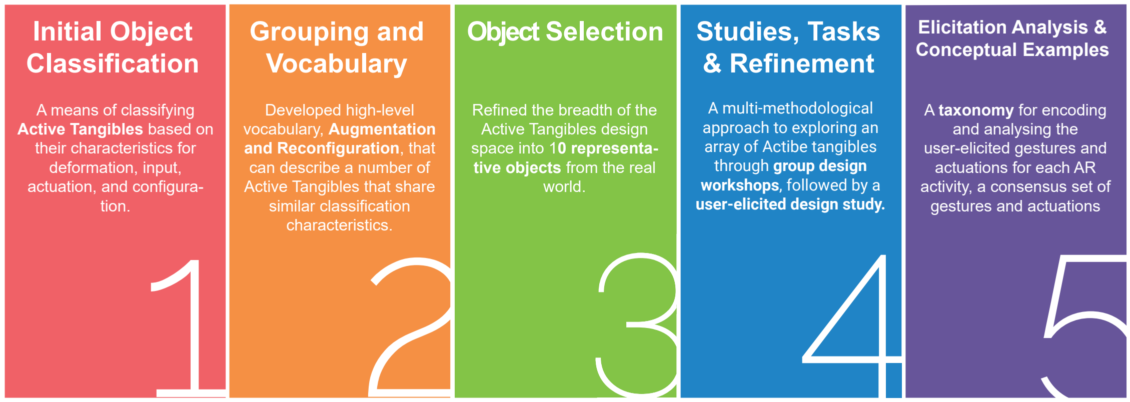

Figure 3.1: An overview of the five different steps taken in our structured approach, each including a summary of their outcome.

2. They make the elicitation results potentially more generalisable to future sensing capabilities by allowing participants to focus on the physical properties of the objects without the influence of sensors.

We aim to gain a deeper understanding of the opportunities for interaction beyond what is currently technically feasible. Our goal is to explore the role of PM-Objects in AR by distilling the novel characteristics of a series of objects to better understand how properties and affordances may be mapped to complex 3D interactions.

To summarise, there are open questions around leveraging PM-Objects for AR applications, particularly how to utilise the novel characteristics of deformation and configurability. This work takes a first step towards elucidating the role of PM-Objects in AR by distilling the novel characteristics of a series of objects to better understand how properties and affordances may be mapped to complex 3D interactions.

3.2 Design Space & Methodology

To understand the potential of PM-Objects as interaction devices for AR, we have developed a structured approach that includes the following steps:

1. Defining a set of characteristics to classify PM-Objects systematically.

2. Creating a vocabulary to group PM-Objects based on common characteristics.

3. Identifying a range of PM-Objects that cover a variety of characteristics to be examined through empirical study around a set of AR tasks.

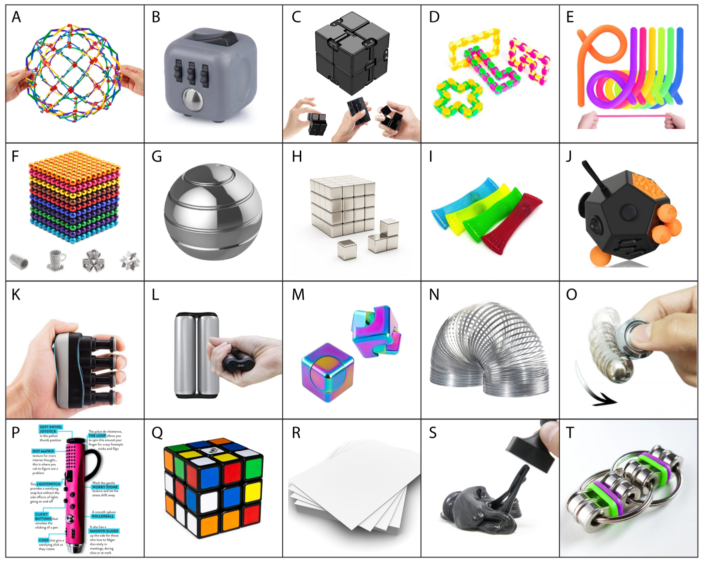

Figure 3.2: 20 objects used to explore physically-modifiable objects: A) Hobermansphere [172], B) Fidget Cube [123], C) Infinity Cube [187], D) Wacky-Tracks [402], E) Stretchy String [360], F) Magnetic Balls [248], G) Gyro Toy [147], H) Magnetic Cubes [249], I) Marble-in-Mesh [254], J) 12-sided Fidget [1], K) Finger Trainer [124], L) ONO Roller [294], M) Cube Spinner [93], N) Slinky [348], O) Magnetic Spinner [250], P)Fiddle Stick [122], Q) Rubik's Cube [328], R) Paper [300], S) Magnetic Putty [380], T) ThinkLink [379].

4. Further refine the PM-Object selection and AR tasks for a user-elicited design study using an AR HMD.

5. Developing a taxonomy for analyzing user-elicited gestures and synthesising them in conceptual examples.

Our work provides a novel means of classifying PM-Objects and uses a combination of existing methodologies to probe them as interaction devices for AR. In this section, we describe the 5 distinct steps of our methodology to distill the properties and affordances of PM-Objects and map them to interaction in AR (see Figure 3.1).

3.2.1 STEP 1: Initial Object Classification

At the outset, we started with 20 exemplar non-digital PM-Objects, shown in Figure 3.2, and made the first iteration on a generalisable framework to describe and categorise them. These were a collection of toys, gadgets, and fidgets that fit the definition of PM-Objects, i.e. objects that can be deformed, reconfigured, and/or have different mechanisms for input. This approach enabled us to explore a breadth of PM-Objects and investigate tangible interaction for AR utilising the recognisable, robust, accessible, and playful nature of the objects. The 20 initial PM-Objects are not exhaustive but represent a breadth of different properties and affordances.

An inductive, iterative thematic analysis [57] concerning the atomic, relational, and spatial properties of each of the 20 objects was conducted by one researcher. This was further developed into a classification based on *properties* and *affordances* (see Figure 3.3). The classification was then reviewed and deductively tested by three independent researchers using a visual worksheet, examples of which are shown in Figure 3.4. While a multitude of methods for classifying physical objects already exists [157, 225, 290], our classification method is unique in describing how physical objects, specifically PM-Objects, can be interacted with and the nature of the input they provide. A PM-Object's *properties* (Figure 3.3), refers to its materiality and capacity for input, and is defined by:

Deformation: Refers to the extent a PM-Object can change its shape. *Mechanical* deformation refers to shape change via some built-in mechanism in the object such as a hinge. *Mechanical* deformation can be *granular* - discrete states of shape change, or *non-granular* - more fluid shape change. *Organic* deformation refers to shape change via the object's material properties, i.e. clay or rubber. A *non-deformable* object has no shape-changing properties.

Input Type: Refers to how input is made available on a PM-Objects, divided into *discrete* and *non-discrete* input. *Discrete* input describes an object that has dedicated input mechanisms that are separate from the object itself for example a button, switch, or dial. *Non-discrete* input describes an object whose input is not separated but can provide input via more holistic actuation, such as a Hoberman sphere expanding and contracting. An additional, more supplementary property for measuring a tangible object's capacity for input is *input-bandwidth*, which describes the number of ways in which you can act upon the object. In Figure 3.4, this is shown on the visual worksheets as a purely relative means of comparison between different PM-Objects.

A PM-Object's *affordance* (Figure 3.3) refers to how it is operated or appropriated [281, 283] and is defined by:

Actuation: Actuation is commonly used in literature to refer to computer-actuated tangibles [314, 338], but in this case, we use it to refer to *user-actuation*: how a user can operate or shape-change the PM-Objects. *Expandable* actuation is where a user can expand the size of an object either through *organic* or *mechanical* means. This actuation can be either *directional* expansion, such as

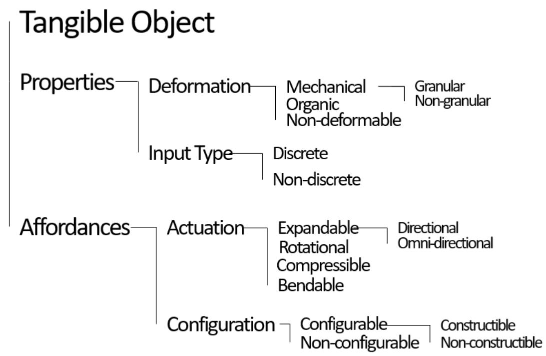

Figure 3.3: The hierarchy for classifying characteristics of PM-Objects based on their properties and affordances. These characteristics are also applicable to tangibles in general and can be used to distill useful characteristics for input in AR.

an elastic band, or *omni-directional* expansion, such as a balloon. *Rotational* actuation is where a user can rotate a particular part of the object but not the whole object itself, for example, a dial or the side of a Rubik's cube. *Compressible* actuation is where a user can partially or wholly compress an object, via *organic* or *mechanical* means, for example, a button press or squeezing clay. *Bendable* actuation is where a user can deform an object's shape, via *organic* or *mechanical* means, for example bending a slinky or an object with a hinge.

Configuration: Refers to the extent to which a user can re-purpose the PM-Objects with some permanence either via construction/destruction or deformation. An example of a *configurable constructible* object is Lego or magnetic cubes. An example of a *configurable non-constructible* object is a Rubik's Cube. *Non-configurable* objects cannot be *constructed* or *destructed* into something new or have their structure changed with permanence.

Outcome of STEP 1: A means of classifying PM-Objects based on their characteristics for deformation, input, actuation, and configuration.

3.2.2 STEP 2: Grouping and Vocabulary

From the classification of tangible objects, we took a further inductive step to develop a vocabulary to group and describe the affordances and properties of a set of physical objects more generally. We define *augmentation* and *reconfiguration* as the general representation of physical object features that can facilitate new interaction concepts when applied to 3D interaction. We postulate that *augmentation* and *reconfiguration*

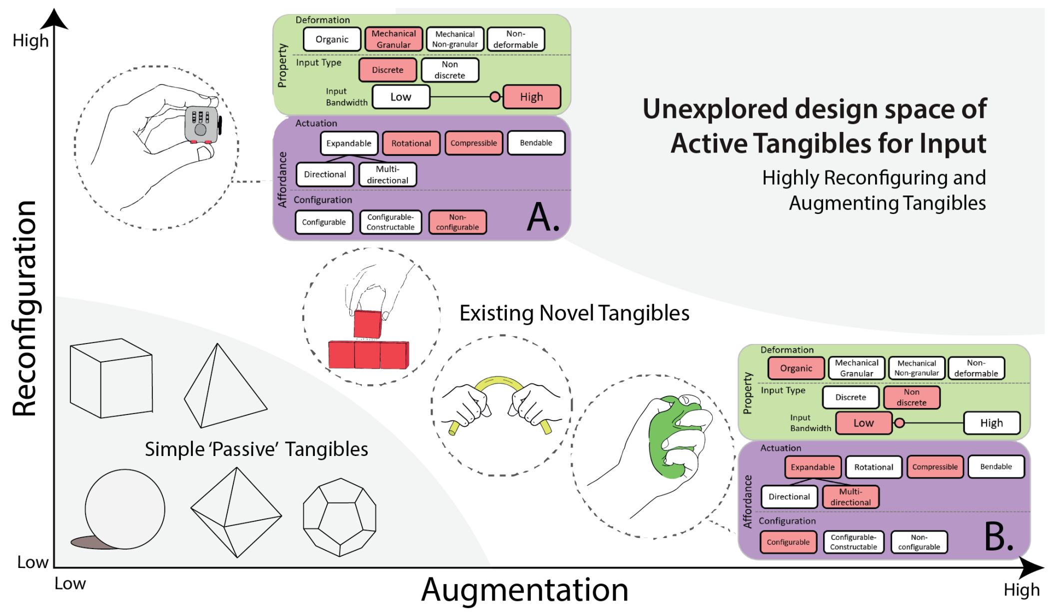

Figure 3.4: A diagram showing some abstract PM-Objects and their arrangement between the concepts of *reconfiguration* and *augmentation*. Included are two exemplar visual worksheets used to review the initial classification (A and B). Highlighted red boxes demonstrate which attributes from the hierarchy belong to which object (see Figure 3.3). Input Heterogeneity denotes the level of diversity in the types of ways an object can be interacted with. These attributes determine the relative position of objects between *reconfiguration* and *augmentation*. Highly augmenting and highly reconfiguring PM-Objects remains an unexplored space.

are two axes for classifying, grouping, and assessing PM-Objects' applicability to 3D interaction, with their position amidst these concepts communicating their general properties and affordances. Importantly, these are not discrete or mutually exclusive categories, rather these are concepts that physical objects might align more or less with. *Augmentation* and *reconfiguration* represent two overarching concepts that embody different features presented in Figure 3.3, and we further define them as:

***Augmentation*:** A tangible object can be considered augmenting if the object's actuation facilitates a unimanual or bimanual gesture that could be otherwise completed unadorned. The main purpose of augmenting objects is to provide motor-control feedback and fine-grain control during gestures that otherwise completed unadorned would be imprecise. Additionally, the object is typically highly malleable with a generally non-discrete, low-bandwidth input.

***Reconfiguration*:** A tangible object can be considered reconfiguring if the actuation of the object provides some means for granular or discrete input, i.e., a button press or dial rotation, with additional support for reconfiguration either by construction, deconstruction, or deformation. The main purpose of reconfiguring objects is to provide multi-modal input for clutch-based interactions, variable control, or manipulation refinement. These objects are desirably configurable for modal interactions, AR system control, or ad-hoc repurposing for different 3D interaction tasks.

However, commonalities exist between the two concepts specifically regarding spatial permanence and passive haptics. A similar inductive process with three researchers was taken with the selected set of PM-Objects to place them between the *augmentation* and *reconfiguration* axes, a subset of which is shown in Figure 3.4.

Following Figure 3.4, a clear unexplored design space becomes apparent. Examples of objects with both low augmentation and low reconfiguration are simple geometric shapes such as cubes and spheres that have already been explored for a variety of different interactions and output mediums for Mixed Reality [111, 235, 310]. Our selected objects are highly augmenting **or** highly reconfiguring and exist on the fringe of what has been explored in related work for AR. Similarly, we could not find any examples of objects that have both high augmentation **and** high reconfiguration in our initial survey of pre-existing objects. Subsequently, we aim to provide an initial characterisation of high augmenting **and** high reconfiguring PM-Objects by exploring objects that exist on the fringe of this design space. Additionally, we also aim to explore the juncture between *augmentation* and *reconfiguration* and identify which properties and affordances are more applicable for different 3D interactions in AR.

Outcome of STEP 2: The outcome of this step directly addresses **RQ1.1**, providing high-level vocabulary – *Augmentation* and *Reconfiguration* – that can describe a number of PM-Objects that share similar classification characteristics. This provides a mechanism to compare and contrast a collection of PM-Objects and deduce the utility of different properties and affordances more generally.

3.2.3 STEP 3: Object selection for Studies

From the initial 20 PM-Objects studied, 10 were selected as a representative sample that covers the range of *augmentation* and *reconfiguration* as shown in Figure 3.5. These objects were chosen to cover the different aspects of the classification hierarchy (Figure 3.3) based on their innate properties for deformation and input, and affordances for actuation and configuration (table in Figure 3.5). Our goal was to identify a diverse selection with as few objects as possible to avoid over-saturation of the object pool and allow ample exploration of each object through empirical study.

O1. **Fidget Cube:** Fidget toy with various dials, switches, buttons, wheels, and joysticks on each face - Figure 3.5a.

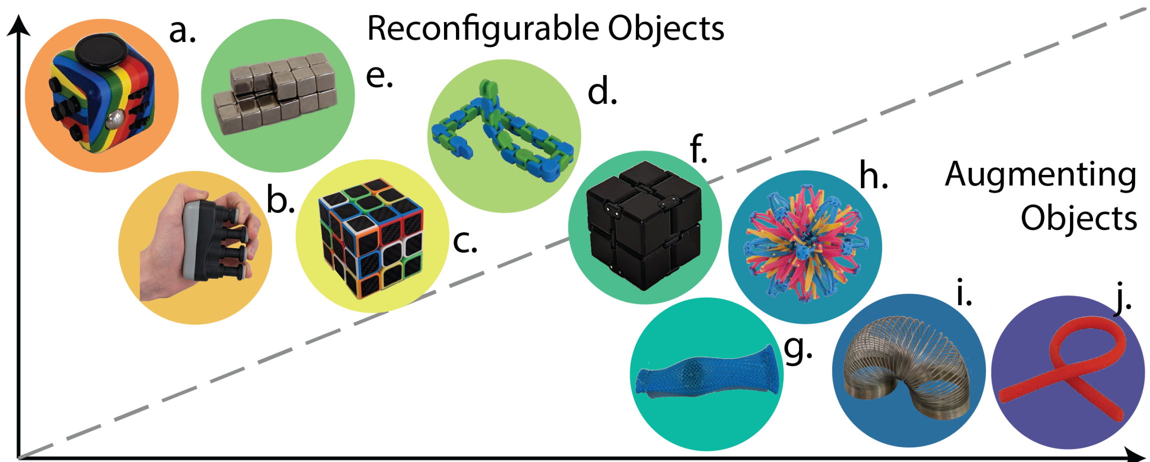

	Object	Fidget Cube (a)	Finger Trainer (b)	Rubik's Cube (c)	Wacky Tracks (d)	Magnetic Cubes (e)	Infinity Cube (f)	Marble in Mesh (g)	Hoberman Sphere (h)	Slinky (i)	Stretch Toy (j)
Property	Deformation	Non-deformable	Non-deformable	Mechanical Non-granular	Mechanical Granular	Non-deformable	Mechanical Non-granular	Organic	Mechanical Non-granular	Organic	Organic
	Input Type	Discrete	Discrete	Discrete	Non-discrete	Non-discrete	Non-discrete	Non-discrete	Non-discrete	Non-discrete	Non-discrete
Affordance	Actuation	Rotational Compressible	Compressible	Rotational	Bendable	Rotational	Bendable	Expandable directional Rotational Compressible Bendable	Expandable Omni-directional	Expandible directional Rotational Bendable	Expandable directional Rotational Compressible
	Configuration	Non-configurable	Non-configurable	Configurable Constructable	Configurable Constructable	Configurable Constructable	Configurable Non-constructable	Configurable Non-constructable	Non-configurable	Non-configurable	Non-configurable

Figure 3.5: An overview diagram of the 10 representative objects selected for our study. The diagonal division shows the arbitrary grouping that can be made between objects that better facilitate reconfiguration (top left) or augmentation (bottom right). The table provides more details on the classification of properties and affordances in each object and demonstrates how each object is a unique combination of different types of deformation, input, actuation, and configuration.

O2. **Finger Trainer:** Grip/finger strength trainer with low resistance pressure switches for each finger - Figure 3.5b.

O3. **Rubik's Cube:** A classic puzzle cube with rotating sides and various coloured faces - Figure 3.5c.

O4. **Wacky-Tracks:** Sensory chain toy, with hinged links that can be assembled into rigid structures - Figure 3.5d.

O5. **Magnetic Cubes:** A series of small magnetic cubes that can be assembled in various configurations - Figure 3.5e.

O6. **Infinity Cube:** A fidget toy composed of 8 hinged cubes, allows for 'infinite' folding and unfolding - Figure 3.5f.

O7. **Marble-in-Mesh:** A fidget toy of nylon mesh with a marble inside that can be slid back and forth - Figure 3.5g.

O8. **Hoberman Sphere:** A mechanical ball that can be expanded and contracted by pulling/pushing on either side. It uses a series of small links that fold and unfold in a scissor-like motion - Figure 3.5h.

O9. **Slinky:** A large springy metal coiled helix. A classic children's toy that can stretch and reform itself - Figure 3.5i.

O10. **Stretch Toy:** A durable string sensory toy designed to be stretched, squeezed, and pulled - Figure 3.5j.

It is important to note that while there are overlaps in terms of these object form factors (several cubic objects), material (mostly plastic objects), and size (mostly palm-sized) the studied objects are explicitly varied on their characteristics for input described in Figure 3.3. The table in Figure 3.5 demonstrates how the studied objects embody these different characteristics. The positioning of these objects in Figure 3.5, signifies a loose grouping that can be made depending on whether the object better facilitates *reconfiguration* or *augmentation*. The purpose of grouping is not to provide absolute labelling of the physical objects but to allow for a high-level, thematic discussion of multiple objects that are related based on *properties* and *affordances*.

Outcome of STEP 3: We have refined the breadth of the PM-Objects design space into 10 representative objects from the real world. This set covers a wide range of characteristics from *augmentation* to *reconfiguration* and will be used to further probe the design space of PM-Objects in AR.

3.2.4 STEP 4: Studies, Tasks, and Object Refinement

To explore this refined set of PM-Objects we conducted two empirical studies: a number of **group design sessions** with 2 to 3 participants each performed outside of AR, followed by an individual **user-elicited design study** in AR.

For the first study, we used the entire set of 10 objects listed previously. The study was performed outside of AR with a series of AR manipulations simulated on a 2D display, ranging from virtual object selection, transformation, modelling, and editing. Inspired by the Research through Design (RtD) method of 'card-sorting' [295] and guessability methodology [308, 418, 419], participants were asked to collectively explore the physical objects, determine how they would perform the task with each object, and then rank the objects from the worst to best for that task. We found that participants described *reconfiguring* and *augmenting* objects differently, and that while *reconfigurable* objects were generally ranked higher, *augmenting* objects were ranked higher for spatially complex modelling tasks. Refer to Section 3.3 for more details on the group design workshops.

The second study was performed using an AR HMD and focused on a refined set of tasks, which were projected for participants in AR following a common guessability methodology [308, 419]. The tasks for this study followed the same categories as the previous study, but we purely focused on virtual object interactions: **transformation**, **modelling**, and **editing**. To avoid participant fatigue and strain, we repeated each task with only the top 4 ranked objects of the previous study for that category. We found that while generally *reconfigurable* objects were preferred for more tasks, while gestures using *augmenting* objects were typically much more agreed upon. Refer to Section 3.4 for more details on the elicitation study.

Outcome of STEP 4: A multi-methodological approach to exploring an array of PM-Objects properties and affordances for different canonical AR activities. A group design session to openly explore the objects, followed by a user-elicited design study to produce a set of gestures and actuation using the set of PM-Objects.

3.2.5 STEP 5: Elicitation Analysis and Conceptual Examples

To encode and analyse the results of the user-elicited design study, we developed a taxonomy of tangible object gestures and actuation. We expanded on preexisting hand gesture taxonomies for AR/VR elicitation studies to include dimensions that codify different types of actuation, configurations, and spatial mappings to virtual objects [70, 215, 308, 418, 419]. Following the process of encoding, we developed a consensus set of tangible object gestures and actuation. The consensus set is a visual depiction of the most agreed-upon gestures across the set of physical objects and AR tasks, resulting in the most common types of actuation, configurations, and spatial mappings to the virtual objects.

Furthermore, we operationalised the results from the consensus set and participant discussion from the group design sessions, into conceptual examples of PM-Tools. These examples are not an exhaustive representation of the results, but rather an example of how we extrapolate commonly utilised properties and affordances from user elicitation, and combine them into a new form factor for 3D interaction in AR.

Outcome of STEP 5: A taxonomy for encoding and analysing the user-elicited gestures and actuations for each AR activity, a consensus set of gestures and actuations, and two conceptual examples of PM-Objects interaction devices based on our results.

3.3 Study 1: Group Design Workshops

The group study involved the 10 different PM-Objects discussed previously and an array of AR tasks presented using 3D animations on a 2D display (see Figure 3.6). To reiterate, the study was not focused on the PM-Objects themselves, but rather the characteristics that they embodied. Note that while AR tasks were demonstrated on a 2D display, participants were explicitly instructed to imagine as if the virtual

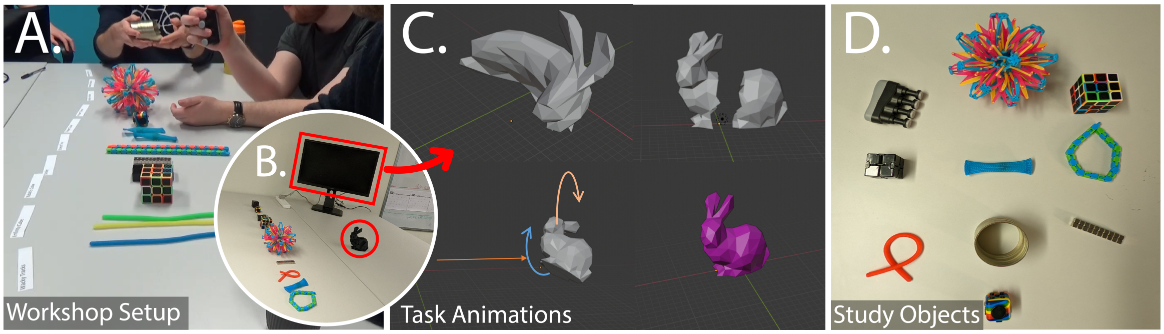

Figure 3.6: The study setup for the group design workshop (A). Tasks were demonstrated using fish-tank VR on a 2D display (B and C). Our 10 study objects that participants explored are also shown.

model was within their physical space. AR-mediated by HMD was not used at this stage, as we did not want to hinder participant communication. The group study was primarily used to present examples of different AR tasks varying in complexity across four topics (shown in Table 3.1), to compare the relative affordances of all PM-Objects, and openly discuss how they might be used to complete each AR task. In this study, we aimed to investigate **RQ1.2** and understand which groups of PM-Objects, between *augmentation* and *reconfiguration*, are preferred for which AR tasks and divulge collective participant reasoning.

3.3.1 Procedure, Tasks and Participants

For each task, participants were asked to collectively rank the 10 PM-Objects from first to last based on how appropriate the object would be for the task. Participants were encouraged to 'think aloud' and discuss with one another the advantages and disadvantages of each object for a given task. Participants were asked to ignore the issues of system tracking or recognition of the objects to allow for free thinking during the group design session.

To reiterate, the 10 tangible objects selected for the group study were curated by a rigorous design space analysis in line with the classification discussed in Section 3.2 (see Figure 3.3 and 3.5). We aimed to cover the design space with as few objects to avoid over-saturating the object pool and make object ranking feasible for participants. For both the group study and elicitation study, we used solely 3D interaction tasks with a distinct focus on virtual object selection, transformation, modelling, and editing. After reviewing common topics and tasks in previous research [79, 230, 296, 308, 384, 400, 401, 419], we devised sixteen tasks in total for the group study spanning four topics: Selection, Transformation, Modelling, and Editing (see Table 3.1). Although previous work has often separated tasks by degrees of freedom,

Topic	Task	Inspired By
Selection	1. Single-Selection 2. Multiple Selection 3. All Selection 4. Partial Selection	Wobbrock et al. [419]
Transformation	5. Translation (all axes) 6. Rotation (all axes) 7. Scale (all axes) 8. Mixed Transformation	Watanabe et al. [405] Piumsomboon et al. [308]
Editing	9. Colour Change 10. Creation, Duplication, & Deletion 11. Grouping & Ungrouping 12. Undo & Redo	Wobbrock et al. [419] Piumsomboon et al. [308]
Modelling	13. Smoothing & Coarsening 14. Extrude & Intrude Face 15. Splitting & Joining 16. Bending, Twisting, & Tapering	Watanabe et al. [405]

Table 3.1: The list of sixteen design workshop tasks for virtual object-centric manipulation, spanning four topics.

such as translation, we instead combined degrees of freedom to increase the complexity of each manipulation.

The chosen virtual model used to convey the AR manipulations was a low-poly model of the Stanford bunny [356] which was a relatively simple model that was non-linear and non-abstract. Choosing a model that is recognisable with easily identifiable features (ears, feet, tail, etc.) to visualise the 3D transformation and modelling manipulations is especially important for elicitation as it aids participant understanding of how manipulation is occurring. For example, abstract 3D models like spheres or cubes do not necessarily have a 'correct' orientation when considering transformation or have easily identifiable features when considering mesh manipulations.

In total, 10 participants were recruited for the study, 9 identified as male and 1 as female. 5 participants were ages 18-24, 4 participants were ages 25-34, and 1 participant was age 35-44. 7 participants had used AR/VR occasionally, 1 participant daily, 1 participant weekly, and 1 participant never. All 10 participants were right-handed. The group study was conducted in a quiet meeting room and participants were positioned around a table with the 10 PM-Objects, each labelled, and a large display for showing each of the AR tasks. The space was audio and video recorded from two different perspectives.

Participants were given a brief introduction to AR and AR interaction before being introduced to the 10 PM-Objects which were all presented simultaneously on a large table. After a 5-minute orientation period with all objects, participants were given tasks in order of complexity: Selection, Transformation, Editing, and Modelling. Inspired by the guessability framework [418] participants were shown the effect of a manipulation on-screen and then asked to think-aloud and discuss how they would use each tangible object to perform that manipulation. Participants were allowed to use the objects in mid-air and/or on a physical surface to complete the manipulation. They would then collectively rank the tangible objects from worst to best for that task. In total, there were 4 group studies, 2 groups of 3 participants and 2 groups of 2 participants, with each session taking approximately 2 hours.

Across all 4 workshops, we recorded the collective participant rankings for each object across the 16 tasks spanning the 4 task topics, for a total of 64 unique rankings (16 tasks x 4 workshops). The results are shown in Figure 3.7, with the task numbers mapped directly to the tasks described in Table 3.1. Audio and video recordings from two perspectives were used to capture the interactions, gestures, discussion, and justification around the 16 ranking tasks for each workshop.

3.3.2 Study 1 Results and Findings

To reiterate, we consider a PM-Object to be either reconfigurable or augmenting if it has more classification characteristics of one over the other (see Figure 3.5). This is not a strict labelling, as some PM-Objects fall in between or have attributes of both, but using these concepts allows us to generalise many properties and affordances simultaneously. Across the 4 task groups, there were general differences between how participants justified and conceptualised *reconfigurable* and *augmenting* PM-Objects depending on the task. Herein, we discuss the main participant observations, the outcome of the ranking tasks, and the implications for the use of PM-Objects.

3.3.2.1 Selection

Reconfigurable PM-Objects were typically ranked higher ($\bar{x} = 17.6$, $\sigma = 7.77$) than augmenting PM-Objects ($\bar{x} = 26.4$, $\sigma = 7.38$) when considering virtual object selection. In general, participants preferred having discrete input on an object, such as the buttons on the fidget cube ($\bar{x} = 12$, $\sigma = 1.63$) or finger trainer ($\bar{x} = 14.25$, $\sigma = 1.71$), as means of either 'cycling' through virtual objects as a form of disambiguation or as means of selection confirmation. Additionally, a 'pointer' metaphor was frequently prescribed to reconfigurable PM-Objects as an alternative disambiguation mechanism, as opposed to 'cycling', due to their resemblance to typical hand controllers. Some reconfigurable PM-Objects were ranked favourably, without having discrete input, as they afforded pointing as an extension of the hand. A 'wand' metaphor was commonly employed by participants as a means of selecting

Task Group / Sub Task

		Selection				Transformation				Editing				Modelling			
	Object	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Reconfiguration	Fidget Cube	12	14	10	12	8	9	4	6	6	6	11	6	4	5	17	16
	Finger Trainer	14	16	12	15	17	22	23	26	14	14	18	9	13	12	25	31
	Rubik's Cube	31	35	33	27	12	4	22	17	14	17	35	11	12	12	35	30
	Wacky Tracks	12	11	19	11	25	21	32	29	21	13	13	18	24	26	13	18
	Magnetic Cubes	18	14	21	15	25	25	35	29	24	17	12	24	19	21	5	23
Augmentation	Infinity Cube	37	37	37	40	35	29	34	36	39	31	32	28	40	40	29	36
	Marble in Mesh	21	25	27	24	29	32	22	20	14	30	31	24	18	18	24	14
	Hoberman Sphere	17	18	17	22	26	17	15	7	26	37	22	33	24	27	29	35
	Slinky	28	20	15	28	17	23	14	20	29	27	23	32	30	26	21	9
	Stretch Toy	30	30	29	26	26	38	19	30	33	28	23	35	36	33	22	8

Rank: 4 – 40

Figure 3.7: The object rankings as a result of the workshop study, are displayed across all tasks. Objects are listed from reconfiguration to augmentation (grouping on the left). Numbers in the table denote the **sum** of all participant rankings of an object for a given task across the four workshops. Hence a value of 4 would mean the object was ranked best across all 4 workshops for that task, likewise, a value of 40 would mean the object was ranked the worst across all 4 workshops for that task. Green denotes a higher ranking (lower number) amongst participants and red denotes a lower ranking (higher number).

through gesture. On the whole, reconfigurable PM-Objects were typically used for selection at a distance, as opposed to selection through direct contact with a virtual object. However, there were some exceptions for augmenting PM-Objects that could be volumetrically deformed/expanded such as the Hoberman sphere ($\bar{x}$ = 18.5, σ = 2.38) and the Slinky ($\bar{x}$ = 22.75, σ = 6.4). Participants frequently used an 'encapsulation' metaphor as a means of direct selection, i.e. changing the physical object's size to spatially encapsulate one, multiple, or part of a virtual object.

3.3.2.2 Transformation

A common observation from participants was that all the physical objects could be used as a spatial proxy for the virtual object being transformed. In this case, they typically ranked reconfigurable PM-Objects higher ($\bar{x}$ = 19.55, σ = 9.49), than augmenting PM-Objects ($\bar{x}$ = 24.45, σ = 8.41), using them as coarse grain spatial proxies and then leveraging the discrete input, such as buttons or dials, either as a mechanism for refinement or as a gain value modifier for the spatial proxy. For all transformation tasks, the fidget cube ($\bar{x}$ = 6.75, σ = 2.22) was typically ranked in the top two. Some participants noted that, for translation and rotation tasks, using an object as a proxy for prolonged periods may cause fatigue. Instead, they compartmentalised the interaction, mapping different translational or rotational axes to different actuation on a physical object. For example, participants suggested that a Rubik's cube's rotating sides or a finger trainer's buttons could be bound to the x,y,z axes. Augmenting PM-Objects were generally ranked higher for scaling ($\bar{x}$ = 20.8, σ = 8.04) and mixed transformations ($\bar{x}$ = 22.6, σ = 11.08) than they were for translational ($\bar{x}$ = 26.6, σ = 6.5) and rotational ($\bar{x}$ = 27.8, σ = 8.11). Volumetric deforming PM-Objects, such as the Hoberman sphere (n = 15), slinky (n = 14), and the stretch toy (n = 19), were ranked higher for scaling. The Hoberman sphere performed particularly well for mixed transformation (n = 7), participants used a 'bubble' metaphor where the sphere was used as a translational and rotational proxy, with actuation mapped to virtual object scaling.

3.3.2.3 Editing

For virtual object editing tasks, reconfigurable PM-Objects were mostly ranked higher ($\bar{x}$ = 15.15, σ = 7.01) than augmenting PM-Objects ($\bar{x}$ = 28.85, σ = 5.82). The fidget cube, in particular, was typically ranked the best across all editing tasks ($\bar{x}$ = 7.25, σ = 2.5). The high *input heterogeneity* was often described by participants as the best attribute of the fidget cube for editing tasks. Constructible PM-Objects, such as the magnetic cubes or the wacky tracks, were ranked favourably for creation, duplication, and deletion (n=17; n=13) and for grouping and ungrouping (n=12; n=13). A common observation is that participants would combine separate pieces of a physical object together to represent an editing action, such as connecting two magnetic cubes together as to group the virtual objects they represent. For editing a virtual object's colour, the marble in mesh was one of the few augmenting PM-Objects ranked highly (n = 14) due to participants employing a 'slider' metaphor to traverse the colour spectrum.

3.3.2.4 Modelling

For modelling tasks, there was often a divide in the object ranking between tasks that were perceived as parameter manipulation, such as coarsening or smoothing a virtual object or changing part of the mesh's convexity, and tasks that were spatially more complex, such as bending or twisting a virtual objects mesh.

Reconfigurable PM-Objects were often ranked higher for smoothing/coarsening ($\bar{x}$ = 14.4, σ = 7.57) and extruding/intruding ($\bar{x}$ = 15.2, σ = 8.29), and augmenting PM-Objects were often ranked higher for splitting/joining ($\bar{x}$ = 25, σ = 3.81) and bending/twisting/tapering ($\bar{x}$ = 20.4, σ = 13.98). Constructable reconfigurable PM-Objects were an exception and were ranked favourably for splitting/joining ($\bar{x}$ = 9, σ = 5.66). For extruding/intruding, participants often would prescribe a spatial 'clutch' metaphor, to compress or expand a virtual objects' mesh, with reconfigurable PM-Objects often preferred leveraging discrete input to turn the clutch on and off.

3.3.2.5 Summary

An interesting observation is that the infinity cube was ranked low for all tasks ($\bar{x} = 34.38$, $\sigma = 4.98$). This could be because it is somewhat of a hybrid PM-Object between reconfiguration and augmentation to the extent that it was difficult to map its affordances for a given task. Another possibility is that its affordances were too similar to other cube PM-Objects, such as the Rubik's cube or the fidget cube, but the infinity cube lacked a mechanism for discrete input. Collectively, participants described augmenting PM-Objects as being 'intuitive', 'natural', and 'satisfying' while reconfigurable PM-Objects were described as 'precise' and 'efficient'. Overall, considering our **RQ1.2**, reconfigurable PM-Objects were ranked more favourably across all tasks, however, certain augmenting PM-Objects scored quite well in specific tasks. For example, the Hoberman sphere was ranked highly for mixed transformations, and the stretch toy and slinky were ranked highly for bending, twisting, and tapering.

3.4 Study 2: User-Elicited Design Study

Following the results of the workshop, we have partially addressed **RQ1.2**, i.e., we have shortlisted certain PM-Objects for different AR activities in the group workshop by taking the highest-rated PM-Objects for each task (see Figure 3.7). We curated the top 4 PM-Objects for each task topic and used them to elicit user-defined gestures in the follow-up elicitation study, with the intent to answer **RQ1.2** and **RQ1.3**. We compiled the most agreed-upon interactions into a consensus set which we can later draw upon for our conceptual examples and general discussion.

3.4.1 Procedure, Tasks, and Participants

We especially selected tasks in transformation, modelling, and editing as we primarily focus on 3D manipulations. This resulted in 14 tasks using the top 4 PM-Objects for each one, based on the results from the group study (see Table 3.2 for task list). The guessability study was conducted with the **same participants** as the group study but on an individual basis. As in the previous study, all tasks were virtual object-centric, but this time we presented the effect of the task by showing a 3D virtual animation using an AR head-mounted display (HMD) using the same virtual model of the Stanford bunny [356] used in the prior study. To reiterate, this model was specifically chosen to aid understanding of the different types of manipulations and aid the explanation of interactions by having identifiable features (ears, tail, feet, etc). Participants were asked to describe and demonstrate the gestures and actuation they would perform using a given object to complete the manipulation. Once again,

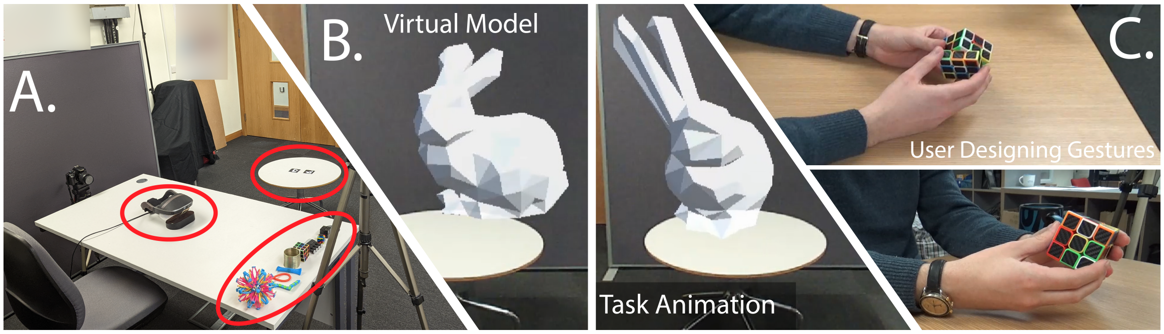

Figure 3.8: The setup for the user-elicited design study. **A)** Shows the Meta AR HMD, the study PM-Objects, and the projection space. **B)** Shows an example AR interaction animation of a virtual model demonstrated in the projection space. **C)** Shows a participant demonstrating a designed gesture for an AR task using one of the PM-Objects.

they were encouraged to follow the think-aloud protocol and ignore issues of object tracking or gesture recognition.

The experimental space was divided into two parts: a table (65×80 cm) which the participants could use to demonstrate a gesture with an object, and a (70×70 cm) circular projection space approximately two meters away from the participant (see Figure 3.8). Each participant was seated at the table whilst wearing the Meta2 AR HMD at 2550×1440 resolution, 60 Hz refresh rate, with a 90°field of view. Two high-definition cameras were mounted to the flat table from two different viewing angles, and the HMD camera captured a participant's gestures. The virtual objects for each task were positioned on the projection surface using a manually positioned virtual marker placed by the experimenter to ensure no occlusion with the environment.

After a short reminder of all the PM-Objects and an introduction to the AR HMD, the experimenter would brief the participant on all 14 tasks as well as which PM-Objects would be used for which tasks. Tasks were presented to the participant in order of complexity, with opportunities for breaks after each task. Although the projection area was outside of a participant's direct reach they could choose either to complete the task directly or indirectly as part of the gesture design process. For each task, participants were given a period to design the chosen gesture and actuation with the given object. Participants could view the animation as many times as they needed and would then perform the final gesture and actuation for the experimenter.

In total, 560 individual gestures were elicited by 10 participants performing 14 tasks with 4 PM-Objects for each task. Video recordings of 3 different perspectives (side-on, top-down, and point-of-view) were analysed: gestures were encoded by 3 independent researchers following the taxonomy defined below, and the qualitative feedback, given by participants following the think-aloud protocol, was transcribed.

Topic	Task
Transformation	1. Translation + Rotation 2. Translation + Scale 3. Rotation + Scale 4. Mixed Transformation
Modelling	5. Smoothing 6. Coarsening 7. Intruding 8. Extruding 9. Splitting 10. Joining 11. Bending 12. Twisting 13. Tapering
Editing	14. Colour Change

Table 3.2: Refined topics and task list for the user-elicited design study.

3.4.2 Tangible Object Gesture and Actuation Taxonomy

Considering previous literature around gesture taxonomies, we were inspired by Wobbrock's original surface taxonomy [419] and Piumsomboom's AR gestures taxonomy [308] which we adapted to encompass interaction with tangible objects for AR. A seven-dimensional taxonomy was developed for object gestures and actuation, using two dimensions from Wobborck's taxonomy: *Form* and *Flow*, and two dimensions from Piumsomboom's taxonomy: *Locale* and *Symmetry*. We omit *Binding* and *Nature* from Wobbrock's taxonomy as all gestures were object-centric, and physically acted on the virtual object. Instead, we devised **three additional dimensions**: *Spatial Binding*, *Actuation*, and *Configuration*. Table 3.3 shows a breakdown of each dimension. Notably, for each studied object certain gestures and/or actuation will not be possible for participants to perform. For example, a Rubik's Cube cannot be bent, or actuated in a symmetric bimanual fashion. So while the taxonomy, by design, covers all possible actions within the pool of studied PM-Objects, each object is not able to have all attributes of the taxonomy.

Form: Similar to Piumsomboom's taxonomy, we omit *one-point touch* and *one-point path* as these are not relevant as gestures and actuation occurs in 3D space. In contrast, for this dimension, we focus on the form of PM-Objects rather than hand gestures which include pose, path, and actuation. The *Flow* dimension consists of *continuous*, when the manipulation happens whilst a gesture or actuation is being performed, and *discrete* when the manipulation happens only

Taxonomy of Tangible Object Gestures and Actuation in AR

Form	Dynamic Pose	Object is pivoted without relocation or actuation
	Static Pose and Actuation	Object is actuated without relocation or pivoting
	Static Pose and Actuation with Path	Object is actuated and relocated without pivoting
	Dynamic Pose and Actuation	Object is actuated and pivoted without relocation
	Dynamic Pose with Path	Object is relocated and pivoted without actuation
	Dynamic Pose and Actuation with Path	Object is actuated, pivoted, and relocated
Symmetry	Dominant Unimanual	Gesture and/or actuation performed only by the
	Nondominant unimanual	Gesture and/or actuation performed only by the nondominant hand
	Symmetric bimanual	Gesture and/or actuation performed by both hands with similar form
	Asymmetric bimanual	Gesture and/or actuation performed by both hands with different form
Locale	On-surface	Object makes contact with physical surface
	Mid-air	Object used in mid-air with no physical contact
	Mixed locales	Object used in both locales
Flow	Discrete	Response occurs after gesture or actuation is completed
	Continuous	Response occurs during gesture or actuation
Spatial Binding	Full proxy	Object is a translational and rotational proxy for a virtual entity
	Translation proxy	Object is a translational proxy for a virtual entity
	Rotation proxy	Object is a rotational proxy for a virtual entity
	Spatial non-proxy	Object is not spatially mapped to a virtual entity but is spatially significant to the interaction
	Semantic	Object is not spatially mapped nor utilises spatiality for interaction
Actuation	Expanded directionally	Object is expanded in a certain direction(s)
	Expanded omni-directionally	Object is expanded in all directions
	Rotating	Object is partially rotated on one or more axes
	Compressing	Object is squeezed or compressed
	Bending	Object is bent or deformed in a particular direction
	No actuation	Object is not actuated
Configuration	Reconfiguration construction	Object is constructed or destructed into a new configuration
	Reconfiguration non-construction	Object is manipulated into a new configuration without construction or destruction
	No reconfiguration	Object is not changed in configuration

Table 3.3: Taxonomy of tangible object gestures and actuation in AR inspired by the taxonomy for Surface [419] and AR gestures [308]. We have used the same category names for *Form, Symmetry, Locale, Flow*, and *Spatial Binding* but in our work, we have altered the definitions to work with PM-Objects gesturing. We also have the addition of *Actuation* and *Configuration* as additional means of encoding PM-Objects gestures.

when a gesture or actuation is completed.

Symmetry: Refers to how the hands behave during a gesture or actuation of a PM-Object. This dimension was introduced by Piumsomboom et al. and we adapt it to include both gesture and actuation of a PM-Object [308]. This dimension consists of *unimanual* performed by the dominant or non-dominant hand, and *bimanual* performed symmetrically or asymmetrically.

Locale: Refers to whether a gesture or actuation was performed in mid-air, on a physical surface, or both. Again, this is adapted from the Piumsomboom taxonomy to include gesture and actuation of a tangible object in the 3 different locales: *mid-air*, *on-surface*, and *mixed-locales* [308].

Spatial Binding: The first dimension we added for gesture and actuation classification, which characterises the spatial relevance a tangible object has on virtual object manipulation. This dimension is divided into five categories *full-proxy*, *translation-proxy*, *rotation-proxy*, *spatial non-proxy*, and *semantic*. A tangible object that is used as a *full-proxy* is mapped one-to-one to a virtual object spatially in all 6 degrees of freedom. A tangible object that is a *translation-proxy* is only spatially mapped to a virtual object's translational degrees of freedom, and if a tangible object is a *rotation-proxy* it is only mapped to a virtual object's rotational degrees of freedom. A *spatial non-proxy* is where the position or rotation of a tangible object is mapped to an attribute of a virtual object that is not spatial, i.e. scale, colour, etc. A *semantic* spatial binding is where a tangible object's spatial movement does not influence interaction.

Actuation: Another additional dimension we have included for gesture and actuation classification. This describes how an object is actuated based on its inherent affordance. The dimension is divided into six categories *expanded directional*, *expanded omnidirectional*, *rotated*, *compressed*, *bent*, and *no actuation*. For example, an object that is *expanded directionally* may be an organically deformable object such as Play-Doh, a slinky, or a stretch toy and is deformed to a particular direction. This can also be achieved mechanically, for example through an object with a telescopic mechanism. Objects that are *expanded omni-directionally* could include, for example, a balloon being inflated or a Hoberman sphere being expanded and contracted. Objects that are actuated via *rotation* can include a mechanical rotation such as a dial being turned or a side of a Rubik's cube being rotated, or organic rotation such as Play-Doh being twisted. Objects that are actuated via *compression* include organic compression through deformation, for example compressing a stress ball, and mechanical compression through a discrete button press. An object can be actuated through *bending* organically, for example, bending a slinky or stretch toy, or mechanically through a joint or hinge. *No actuation* refers to a gesture for manipulation that does not include any actuation of the tangible object.

Configuration: The final dimension we introduce. This describes a tangible object's capacity to be changed into a different arrangement that is significant for interaction either via *construction*, assembling a new object out of components of a pre-existing object (e.g. Lego or magnetic cubes) or *non-construction*, changing the shape of an object without assembling or disassembling a preexisting object (e.g. reconfiguring a Rubik's cube). An object with *no reconfiguration* capabilities means that its arrangement cannot be changed or was not significant for an interaction.

3.4.3 Study 2 Results and Findings

3.4.3.1 Gesture Classification

560 gestures were recorded and classified as shown in Figure 3.9. The most common characteristics within the seven-dimensional taxonomy across all tangible objects were *static pose & actuation*, *asymmetric bimanual*, *mid air*, *continuous*, *semantic*, *rotated*, and *non-reconfiguration*. For the *Form* dimension: 52.7% of all gestures were *static pose and actuation*, 23.38% *dynamic pose, actuation, and path*, 8.45% *static pose and actuation with path*, 8.27% *dynamic pose with path*, 5.94% *dynamic pose and actuation*, 0.72% *static pose with path*, and 0.54% *dynamic pose*. Gestures performed in the **transformation** tasks were predominantly *static pose and actuation* (29.38%) and *static pose and actuation with path* (28.75%). **Smoothing/coarsening** tasks were 100% *static pose and actuation*, **intruding/extruding** tasks mostly *dynamic pose and actuation* (76.25%), **splitting/joining** tasks typically either *dynamic pose and actuation with path* (37.5%) or *dynamic pose and actuation* (35%), **bend/twist/taper** tasks mostly *static pose and actuation* (87.07%), and **colour change** was 100% *static pose and actuation*.

For *Symmetry*: 65.83% of all gestures were *asymmetric bimanual*, 17.27% *symmetric bimanual*, 16.19% *dominant unimanual*, and 0.72% *nondominant unimanual*. Gestures performed in the **transformation** tasks were mostly *asymmetric bimanual* (60%), **smoothing/coarsening** tasks were mostly *asymmetric bimanual* (75%), **intruding/extruding** tasks were mostly either *asymmetric bimanual* (58.75%) or *dominant unimanual* (41.25%), **splitting/joining** tasks were mostly either *symmetric bimanual* (47.5%) or *asymmetric bimanual* (42.5%), **bend/twist/taper** tasks were mostly *asymmetric bimanual* (87.93%), and **colour change** was mostly *asymmetric bimanual* (67.5%).

For *Locale*, 92.63% of all gestures were performed *midair*, 5.22% *mixed locales*, and 2.16% *on-surface*. Gestures performed for all tasks were mostly *mid-air* locale (at least 87.93% for each type of task).

For *Flow*, 93.71% of all gestures were *continuous* and 6.29% were *discrete*. Gestures performed in the **transformation**, **smoothing/coarsening**, **intruding/extruding**, **bend/twist/taper**, and **colour change** tasks were almost entirely *continuous* (at least 95%). For **splitting/joining** tasks 66.25% of gestures were *continuous* and 33.75% were *discrete*.

For *Spatial Binding*, 38.67% of all gestures were *semantic*, 26.98% were *spatial non-proxy*, 12.77% were *spatial rotation-proxy*, 12.23% were *spatial full-proxy*, and 9.35% were *spatial translation-proxy*. Gestures performed for **transformation** tasks were mostly either *spatial full-proxy* (29.38%), *spatial translation-proxy* (27.5%), or

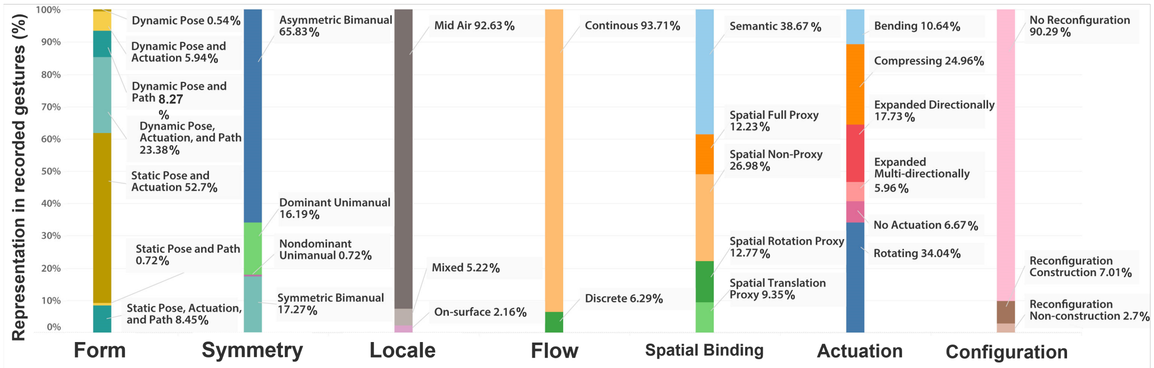

Figure 3.9: Distribution of gesture characteristics, across all participant gestures, based on the tangible object gestures and actuation taxonomy shown in Table 3.3. Each stacked bar represents a category in the taxonomy and shows how different characteristics were distributed across all of the gestures recorded.

semantic (24.38%). **Smoothing/coarsening** tasks were 100% *semantic*, **intruding/extruding** tasks mostly *spatial non-proxy* (71.25%), **splitting/joining** tasks mostly *spatial non-proxy* (58.75%), **bend/twist/taper** tasks were either *spatial translational-proxy* (41.38%) or *semantic* (31.9%), and **colour change** was either *semantic* (70%) or *spatial non-proxy* (30%).

Over 75% of all gestures involved some form of *Actuation* of which 34.04% were *rotated*, 24.96% *compressed*, 17.73% *expanded directionally*, 10.64% *bending*, 6.67% *no actuation*, and 5.96% *expanded omni-directionally*. Actuations performed for **transformation** tasks were mostly *rotated* (39.47%), **smoothing/coarsening** tasks were mostly either *compressed* (50%) or *rotated* (47.62%), **intruding/extruding** tasks were mostly either *compressed* (47.57%) or *rotated* (34.95%), **splitting/joining** tasks were mostly *no actuation* (35.11%), **bend/twist/taper** tasks were mostly either *rotated* (30.65%) or *expanded directionally* (27.42%), and **colour change** was mostly *rotated* (42.86%).

For *configuration*, 90.29% of all gestures involved *no reconfiguration*, 7.01% *reconfiguration construction*, and 2.7% *reconfiguration non-construction*. Gestures performed in the **transformation**, **smoothing/coarsening**, **intruding/extruding**, **bend/twist/taper**, and **colour change** tasks were mostly *no reconfiguration* (at least 87.5% for each task). **Splitting/joining** tasks were mostly either *no reconfiguration* (56.25%) or *reconfiguration construction* (42.5%).

3.4.3.2 Gesture Agreement and Consensus Set

Classifying the 560 recorded gestures resulted in 277 unique gestures across objects and tasks. Following previous elicitation studies, we devised a 'consensus set' which is

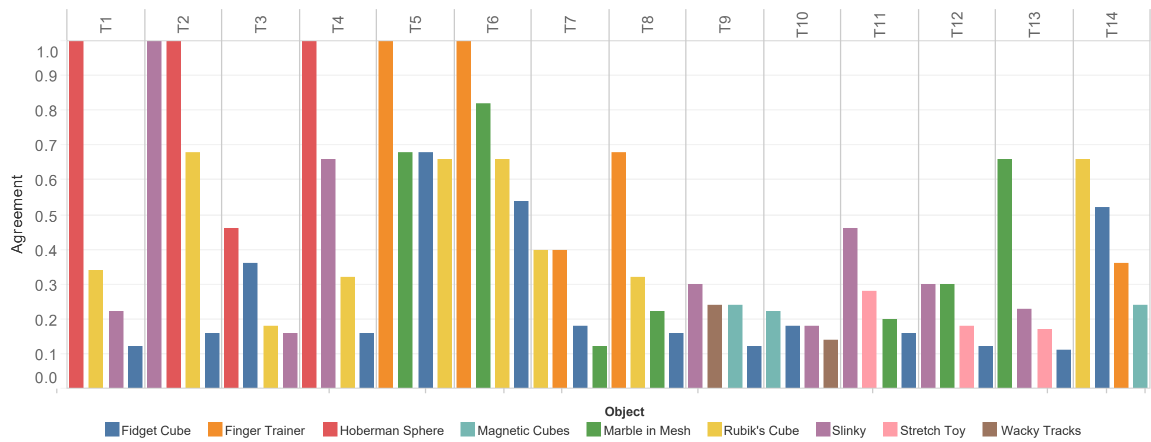

Figure 3.10: Gesture Agreement rates per task, per object in the order of agreement. TX denotes the task, followed by the name of the object used for that task.

comprised of the largest groups of similar gestures that are performed for a particular task [397]. For our study, each unique gesture was scored 1 point for each participant that performed the gesture for a maximum of 10 points. To determine the level of agreement among the defined gestures, an agreement rate was calculated using Equation 3.1 [397]:

$$AR(r) = \frac{P_t}{P_t - 1} \sum_{P_s} \left(\frac{P_s}{P_t} \right)^2 - \frac{1}{P_t - 1} \tag{3.1}$$

Where P_t is the total number of gestures and actuation within the task t, P_s is a subset of P_t containing similar gestures, and the agreement rate $AR(r)$ is between 0 and 1. For example, let us take the virtual object **transformation** task *translation and scale* using the Rubik's cube. This task contained four gestures, with scores of 5, 2, 2, and 1 points, and the tasks agreement rate $AR(t)$ can be determined as follows:

$$AR(t) = \frac{10}{10 - 1} \left(\frac{5}{10} \right)^2 + \left(\frac{2}{10} \right)^2 + \left(\frac{2}{10} \right)^2 + \left(\frac{1}{10} \right)^2 - \frac{1}{10 - 1} = 0.27 \tag{3.2}$$

Figure 3.10 shows the agreement rates for all objects across all tasks. For the consensus set, we used the majority consensus for each task and physical object, i.e. gesture groups with more than 50% agreement rate. This resulted in 18 unique gestures that spanned 8 out of 14 tasks. For the remaining 6 out of 14 tasks, we took the most agreed-upon gestures and included them in the consensus set. The final consensus set contained 25 unique gestures out of the total 277 unique gestures, with the remaining 252 as the discarded set

[397, 419]. The consensus set represented 193 out of 560 gestures (193 out of 560 points, 34.46%). Gestures in the set with majority consensus represented tasks as follows: **transformation** (11.07%), **smoothing/coarsening** (12.14%), **intruding/extruding** (3.57%), **splitting/joining** (1.61%), **bend/twist/taper** (3.39%), and **colour change** (2.68%) for a sum of 35.74%. When considering the agreement rates and the consensus set, it is important to note that some gesture groups were low in agreement but contained gestures with high individual rates.

3.4.3.3 Findings from the Consensus Set

Out of the 25 unique gestures, 13 were using reconfigurable objects, and 12 used augmenting objects. In the consensus set, gestures with reconfigurable objects had an average agreement of 57.9% while augmenting object gestures were 66%. For reconfigurable object gestures in the consensus set, 1 was for **transformation** tasks, 6 for **smoothing/coarsening**, 3 for **intruding/extruding**, 1 for **splitting/joining**, 0 for **bend/twist/taper**, and 2 for **colour change**. Likewise, 6 of the augmenting object gestures were for **transformation** tasks, 1 for **smoothing/coarsening**, 0 for **intruding/extruding**, 1 for **splitting/joining**, 3 for **bend/twist/taper**, and 0 for **colour change**. During the guessability study, participants were encouraged to follow the think-aloud protocol when designing their gestures, and posthoc analysis of the observational recordings showed some common themes and interaction metaphors among participants. Herein, we describe the consensus set per task (shown in Figure 3.11) and address **RQ1.2** and **RQ1.3**.

Transformation: Transformation gestures made up 7 out of 25 in the consensus set with an average agreement of 82.9%. Augmenting gestures were the most commonly agreed on for this task, with most participants adopting a ‘proxy’ metaphor for manipulating the virtual object. For example, the physical object has a 1-to-1 spatial mapping to the virtual object, either directly with the virtual object projected over the physical object or indirectly with the virtual object projected at a distance. This was common for both translation and rotation. However, the scale was often mapped to the physical object’s actuation such as the Hoberman sphere’s *deforming omnidirectional expansion*, the Slinky’s *deforming directional expansion*, or the Rubik’s cube’s *rotational* actuation. Gestures using the Hoberman sphere were the most agreed upon with many participants either using the ‘proxy’ or ‘bubble’ metaphor, i.e. the virtual object(s) being manipulated are spatially positioned inside the Hoberman sphere. While augmenting gestures were common, there was one wildly agreed upon reconfiguring gesture for *scale* and *rotation* using the Rubik’s Cube. Participants typically treated the Rubik’s cube as a *rotational proxy* and used the *rotational* actuation as a dial for scaling the object.

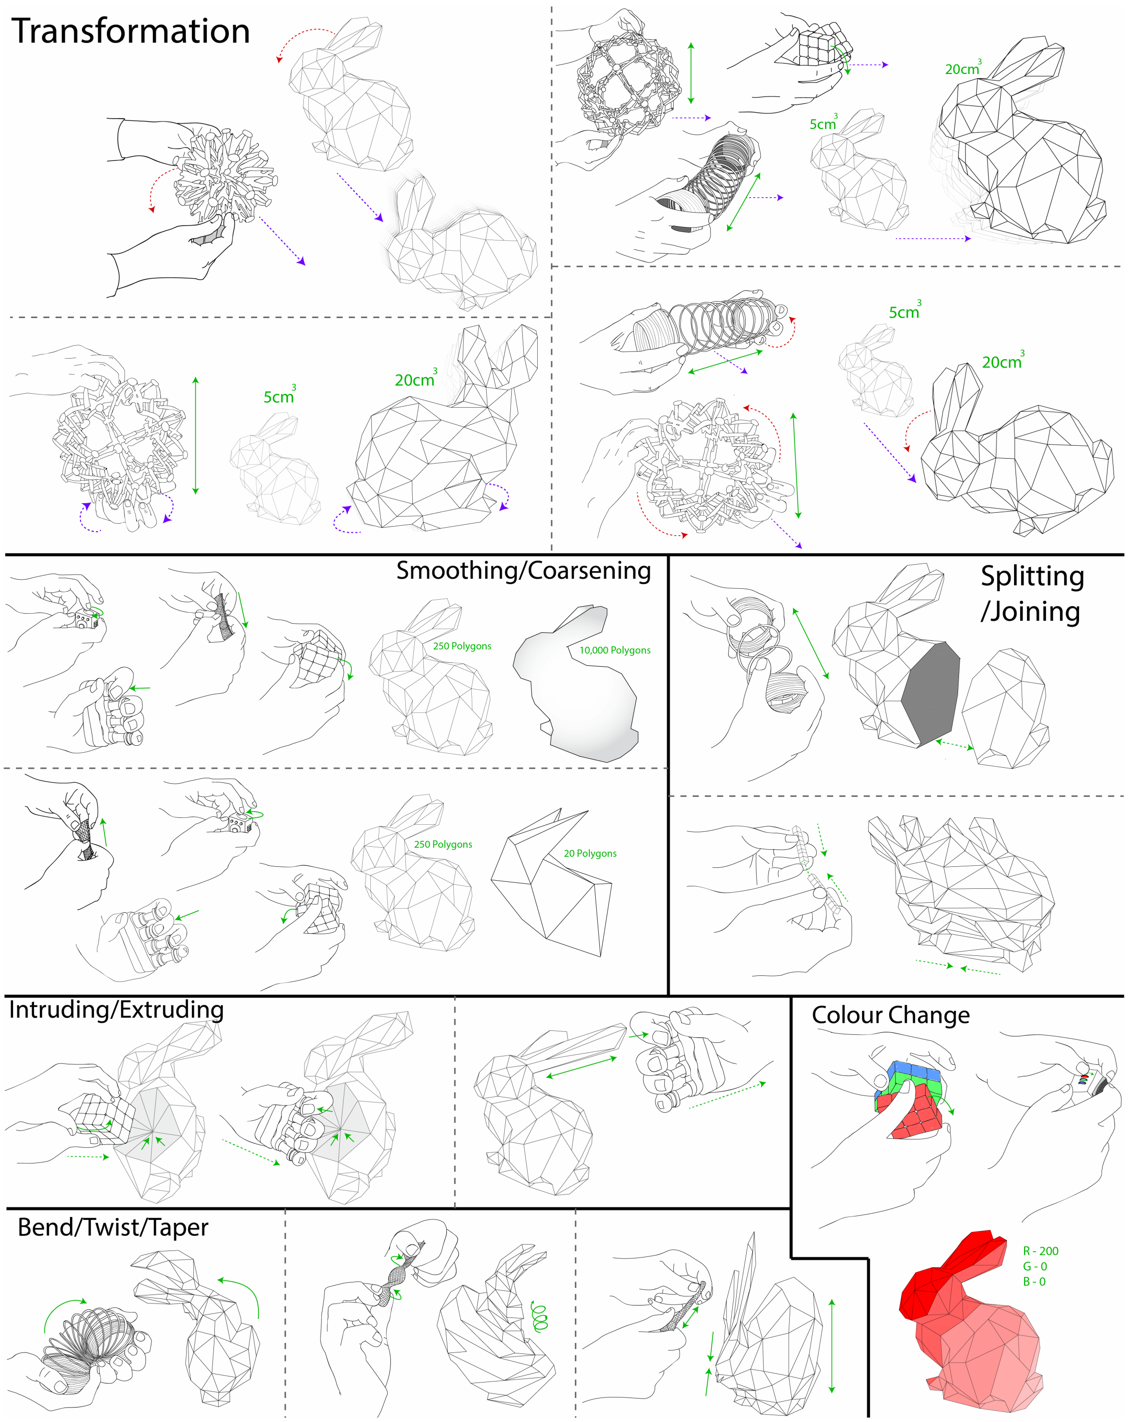

Figure 3.11: The 'consensus set' of gestures (25 in total) from the user-elicited design study separated by task. Dashed arrows convey spatial movement and solid arrows denote PM-Objects actuation or virtual object deformation.

Smoothing/Coarsening: Smoothing/coarsening gestures were the most common in the consensus set, making up 8 out of 25 with an average agreement of 75.5%. Reconfiguring gestures are the most common for this task, except for 2 augmenting gestures using the Marble in mesh. Participants frequently described this task as a 'variable' manipulation task, i.e., changing the polygon count of the virtual object, hence why all gestures are exclusively actuation based without utilising the spatial movement of the physical objects. For the Fidget cube and Rubik's cube, discrete rotational actuation was generally used for both increasing and decreasing the smoothness, while the Finger trainer's buttons were mapped to either increasing or decreasing smoothness. The Marble in mesh, as the only augmenting object for this task, was typically perceived and operated as a 'slider' to manipulate the same variable leveraging a combination of *expansion* and *compression* actuation.

Intruding/Extruding: Intruding/extruding gestures made up 3 of the 25 in the consensus set with an average agreement of 49.3%. Reconfiguring gestures were exclusively used for these tasks, typically using a combination of spatial movement and discrete actuation. For the Finger trainer, a 'clutch' metaphor was popular among participants, where the physical object is positioned near the desired part of the virtual object mesh and a *compression* actuation is applied using the buttons to 'clutch' and then spatially manipulate the mesh to create intrusions or extrusions. By contrast, gestures using the Rubik's cube leveraged the same spatial positioning, for selecting a desired part of the mesh, but instead used a 'dial' metaphor to intrude using *rotational* actuation. Some participants also used this gesture for extruding the mesh it was not the majority consensus.

Splitting/Joining: Splitting/joining gestures made up 2 out of the 25 in the consensus set and was the least agreed upon with only 26% average agreement. One was an augmenting gesture with the Slinky, using a granular *directional* expansion to indicate the position of the split and the degree of separation. For example, a participant can indicate the position of the split along an axis of the virtual object using a combination of the rotation of the Slinky to indicate the axis and then performing an *expansion* on a particular ring to indicate the position of the split. However, some participants indicated that this approach would be constrained when performing non-linear splits such as a concave or convex split. Interestingly, a similar approach was not as popular for joining two virtual objects. The other gesture was a reconfiguring gesture with the Magnetic cubes, using a 'proxy' metaphor to begin with, i.e. one set of magnetic cubes spatially mapped to one virtual object, and then *reconfiguring* the two sets of cubes into one physical object to convey 'joining'. Some participants noted magnetism to be a constraining factor for reconfiguring a physical object due to the polarity of the individual magnets.

Bend/Twist/Taper: Bend/twist/taper gestures made up 3 out of 25 in the consensus set with an average agreement of 47.3%. Augmenting gestures were exclusively used, predominantly leveraging *deformation* actuation to indicate both the degree and direction of the bend, twist, or taper. Gestures with the Slinky were the most agreed upon for the bending task in which participants would treat the Slinky initially as a rotational proxy of the virtual object and then directionally *expand* the Slinky to convey the direction and degree of the bend. Some participants noted that the manipulation was constraining with the Slinky as the hands can only expand the object so far and so a 'clutch' metaphor was described by some to 'turn the object on/off' to allow for user repositioning. For both twist and taper, gestures using the Marble in mesh were the most agreed upon. Similar to the Slinky, the Marble in mesh was initially used as a rotational proxy, to indicate the position of either the twist or taper, however, the position of the marble within the mesh was used to convey the degree and direction of either the twist or taper. For instance, when twisting, the marble position conveys the position or 'limit' of the twist in the virtual object mesh, and the bidirectional rotation deformation of the object indicates the degree of the twist. Similarly, when tapering, the marble position indicates the direction and the *expansion* deformation indicates the degree of the taper. Participants noted a similar issue with the Marble in mesh that the virtual object manipulation was constrained by the degree to which a user/PM-Object could twist or taper. In this case, some participants proposed maintaining a gesture to increase or decrease the degree of mesh manipulation, akin to holding down a button to increase or decrease a value.

Colour Change: Finally, colour change gestures made up 2 out of 25 in the consensus set with an average agreement of 59%. Reconfiguring gestures were exclusively used, applying discrete actuation to manipulate the RGB values of a virtual object. For the Rubik's cube, many participants mapped the RGB colour values (0-255) to 3 rotational 'dials' leveraging either the distinct faces on the cube or the 3x3 sections (as shown in Figure 3.11). Likewise, the Fidget cube had 3 'jog wheels' that many participants used for increasing and decreasing RGB values. Many participants described this task as a matter of 'manipulating a set of variables' comparable to the smoothing/coarsening task. Some participants noted that other colour values such as HSV could be mapped to the same input and then, using a modal approach, interfaced between by utilising an additional actuation of the physical object.

3.5 Conceptual Examples of PM-Tools in AR

While the results from our gesture elicitation study can be interpreted, operationalised, or applied in various ways, we want to demonstrate (by example [226]) how these results can be used. Using the findings from the gesture consensus set and common themes brought up in the design workshop, we ideated the **interaction design** for 2 physically-modifiable tools (PM-Tools) for object-centric AR tasks focusing on *selection and transformation* and *3D Modelling*. For each type of task, we use a combination of the most agreed-upon gestures and popular interaction metaphors from both the group design workshops and the user-elicited design study.

3.5.1 Example 1: The Bubble

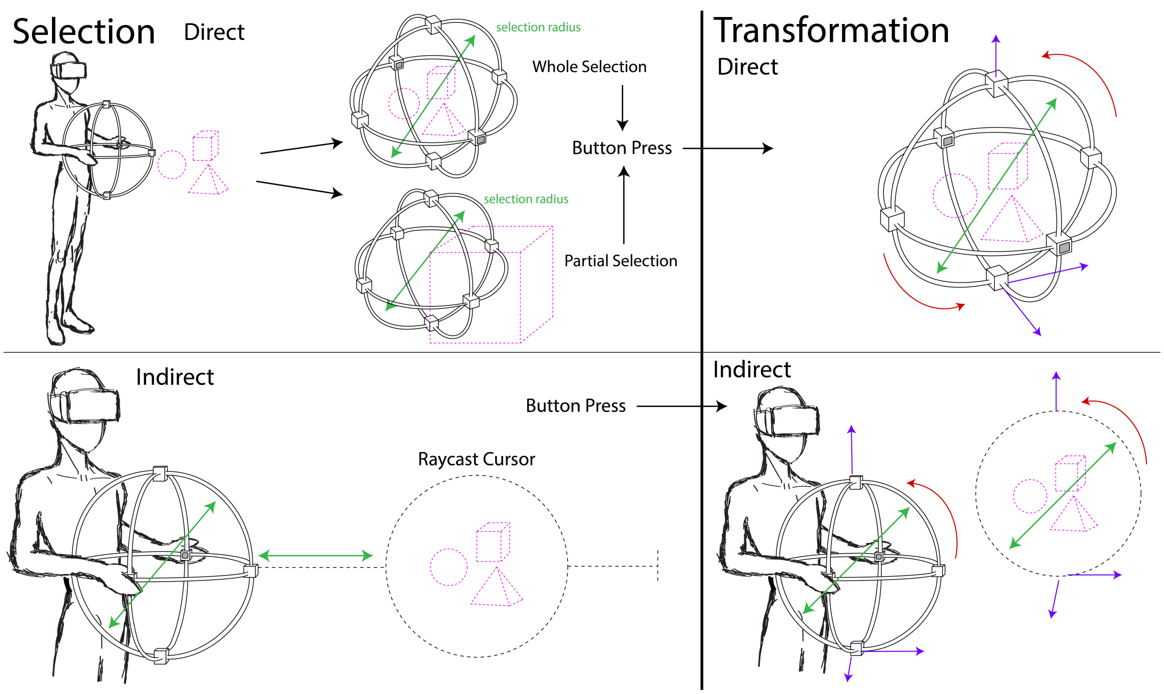

Figure 3.12: 'The Bubble' - A conceptual PM-Tool for partial/whole selection and 6DoF transformations based on the consensus set gestures and common interaction metaphors as expressed by participants.

Based on the consensus set and group study, we conceptualised a PM-Tool for 3D transformation and selection tasks in AR. A common theme from the user-elicited design findings was that participants preferred a mixture of direct and indirect interaction, i.e. transformations with the virtual model(s) collocated with physical transformations at a distance. Hence, the Bubble would be designed for both direct and indirect interactions for transformation and selection. **The Bubble** is primarily based on the *omnidirectional expansion* actuation of the Hoberman sphere, however, we designed some additional desirable properties based on the Fidget Cube – *mechanical granular deformation* and *discrete input* – to enable some common interaction metaphors raised in the study. These additional properties would allow a user to expand the device incrementally so that an exact expansion value from the Bubble may be mapped to a virtual object scale or a selection area. The addition of a dynamic gain value would allow for selection and scaling beyond the constraints of the device. Lastly, *discrete input* could be incorporated to allow for clutching, modal interaction, and control over gain values (see Figure 3.12).

Basic Operation: To perform direct selections, a user simply expands the Bubble to the desired selection area encompassing one or more virtual objects inside and then presses one of the buttons on either of the 'clutch' sections of the object to confirm the selection. The same process can be used for partial object selection, expanding the Bubble to the desired selection area and then placing it over the desired section of the virtual object. For direct transformation, the Bubble utilises a 'clutch' metaphor where a user simply toggles one of the buttons to enable 3DoF manipulation and scaling using the expansion actuation. For indirect selections, the Bubble can be used as a raycast pointing device to select one or more objects at a distance. The object expansion actuation is mapped to the depth of a raycast cursor [28] and by toggling with one of the buttons can be locked in place while the expansion of the object is then mapped to the raycast cursor selection area. A final button press confirms the selection. Partial virtual object selection can also be done in this way. Indirect transformation works in the same way as direct manipulation, once a selection is made, but with the virtual object(s) manipulated at a distance.

Affordance Mapping:

- *Omnidirectional expansion*: Selection area and scaling control.
- *Mechanical granular deformation*: Control over selection area, scaling, and indirect selection of gain values.
- *Discrete input*: Modal input via button press.

3.5.2 Example 2: Sculpting Stick

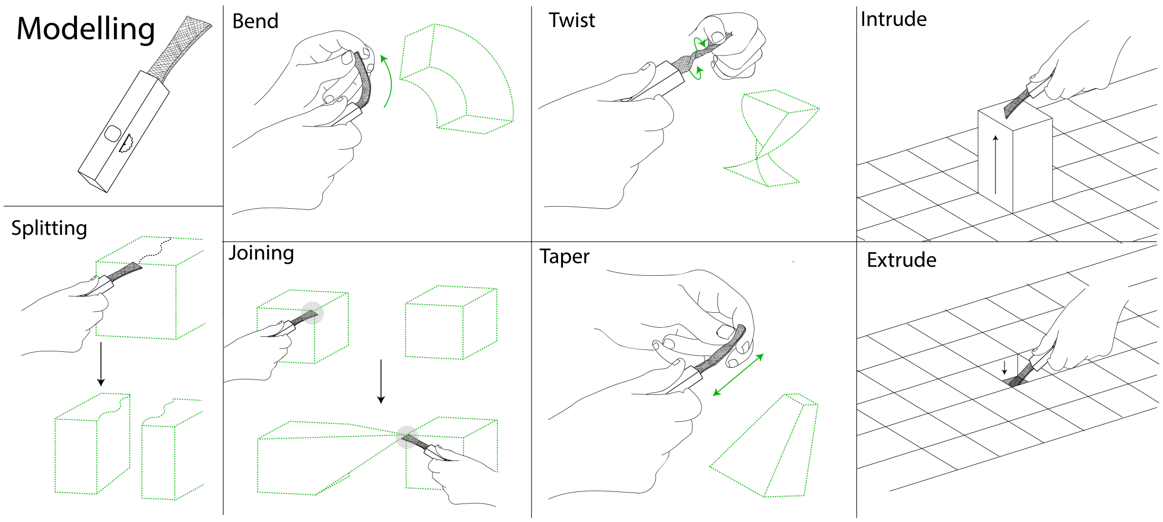

Figure 3.13: 'The Sculpting Stick' - A conceptual PM-Tool for simple 3D modelling based on the consensus set gestures and common interaction metaphors as expressed by participants.

The second tangible device we conceptualised for simple 3D modelling is the Sculpting Stick based on a combination of the *expandable, rotational, compressible, and bendable actuation* of the Marble in mesh, and the *discrete input* properties of the Finger trainer and Fidget Cube (button and dial input). Combining these properties and affordances enables a user to perform simple 3D modelling interactions such as splitting and joining meshes/vertices and performing different types of mesh deformation (see Figure 3.13).

Basic Operation: Using the dial, users can cycle through different modes of the Sculpting Stick. One mode allows a user to manipulate individual vertices of a virtual objects mesh by selecting the vertices directly, using the button to toggle clutching on/off, and then spatially manipulating the vertices. The vertices can then be attached to other virtual object meshes. Similar to vertex manipulation, faces can also be manipulated to allow for concave and convex manipulations of an object's mesh. Another mode of the Sculpting Stick is the knife tool for dividing virtual object meshes using a spatial path. This can again be toggled using the button. Further mesh deformations such as bend, twist, and taper can be achieved by manipulating the flexible 'mesh' tip of the Sculpting Stick with the degree of the deformation increased or decreased via button input.

Affordance Mapping:

- *Organic Deformation and Actuation*: Expandable, rotational, compressible, and bendable input using a flexible tip for organic mesh deformation. The rotational dial allows for modal interaction.
- *Discrete input*: Allows for control and recovery of selection and manipulation, as well as control over gain values via button input.

3.6 Discussion

The central goal of this chapter was to establish a method for *Tool-making* and explore the mediating role of PM-Tools through a systematic exploration of how physical configurations, affordances, and spatial relations map to a range of canonical AR interactions. Herein, we reflect on the methodology, findings on PM-Tools in AR, and limitations of our work.

3.6.1 PM-Tools Framework, Classification, and Insights

PM-Tools have mostly been explored in AR using point examples [17, 36, 220]. However, tangible AR and PM-Tools inherit similar problems from TUI research, namely that transferring the insights from one application domain to another remains fundamentally problematic [45, 129, 199]. To compare TUI concepts in a meaningful way, there is a requirement to lift them into the same frame of reference using similar abstraction, exemplars, and vocabulary. While we can use the insights from TUI literature for tangible AR [125, 126, 191, 189, 338, 387] and while methods for describing tangible AR interaction do exist [46], it is unclear how to derive insights from one point example and apply them in another application domain in a principled manner. This is even more so the case for emerging reconfigurable or augmenting interfaces such as PM-Tools. Additionally, comparing and operationalising different tangibles in AR abstracted from interaction context to focus on affordance is by no means 'straightforward' [178]. Our work provides (i) the vocabulary to describe, generalise, group, and compare a variety of PM-Objects, (ii) a structured methodology for using exemplar PM-Objects that exist in the real world, and (iii) a curiosity-driven approach to exploring these devices for canonical AR tasks.

While we mainly focussed on HMD AR, PM-Tools can be utilised for other mediums of Mixed Reality. Our consensus set of gestures makes use of a variety of direct and indirect interactions for a variety of manipulation and modelling tasks on a virtual object. While direct interaction was always a possibility for participants, indirect gestures are the most common in the consensus set and agreed upon as the most obvious means of achieving a task. The exception here is for intruding and extruding tasks, which were specific to certain areas on a model, suggesting that direct action is preferred when high accuracy is required. Due to the prevalence of indirect gestures, our consensus set gestures and even components of our conceptual exemplars are transferable to other forms of Mixed Reality such as projection AR, fish-tank VR, and HMD VR. Most designed gestures using the study objects were bimanual interactions with a mixture of symmetric and asymmetric gestures. We observed similar behaviour in the bimanual gestures that have been widely documented in related work [63, 144, 167, 169, 170], in which the hands would often adopt different roles when actuating a PM-Tool. Due to the prevalence of bimanual gestures, it's difficult to see how our insights could be applied to handheld AR without exploring more objects that afford unimanual interaction - such as the finger-trainer. Another possibility is to take the properties of augmenting and reconfiguring PM-Tools and apply these to the actual handheld device as peripheral controls that can be leveraged when needed for certain AR interactions.

To summarise the takeaways related to the PM-Tools framework, we introduce and characterise the design space of PM-Tools and explore the possibilities of highly augmenting/reconfigurable interaction devices, applicable both inside and outside of HMD AR. We introduce vocabulary, characteristics, and a means of grouping PM-Objects as well as combining several pre-existing research methods in a novel way to systematically explore these objects for a given interaction space.

3.6.2 Affordance Mapping

A key problem in any emerging technology is to explore, in a nuanced and systematic way, how the action possibilities of the technology match up with wider requirements of the activities and goals of users. Common methodologies, such as guessability studies or gesture elicitation studies [70, 215, 295, 308, 418, 419] and 'card-sorting' [295], are examples of how a mapping between technological capabilities and user goals can be systematically explored. Our work follows a unique combination of these methodologies and addresses the key challenge of how to compare a set of PM-Objects to various fundamental activities within AR. Through the 5-step method described, we can refine our exploration of PM-Objects over time in terms of the study objects and the AR tasks themselves. The outcome of exploring such a mapping is a series of key insights into how various configurations of the *reconfiguration* and *augmentation* design space enables different types of activities in AR. We demonstrate how affordance mapping with preexisting PM-Objects can be a useful, low-cost, and prompt activity for designing affordances in tangible interaction devices as shown in the conceptual examples.

While we acknowledge design guidelines are difficult to achieve, there are some implications from our findings. Firstly, it's important to note that we have two axes for the analysis of affordance mapping: (i) Collective perception of the applicability of affordances for a given task (rankings from Study 1), and (ii) participant consensus on how an interaction should be designed for a given object and task (level of agreement and consensus set from Study 2). Interestingly, the most 'applicable' PM-Objects (objects ranked higher in Study 1) often had the lowest amount of consensus for how the interaction should be designed in Study 2. This was particularly prevalent for more *reconfigurable* PM-Objects.

In general, *reconfigurable* PM-Objects were consistently ranked higher for transformation tasks (translate, rotate, and scaling) and *augmenting* PM-Objects ranked higher for certain modelling tasks (bending, twisting, and tapering), but the design of the interactions was not always agreed-upon for reconfigurable PM-Objects (only one gesture for transformation in the consensus set). This could be due to higher *input heterogeneity* leading to less consensus or that there are more intricacies to consider when designing interaction with reconfigurable objects — the more reconfigurable the more difficult it is to devise a commonly agreed-upon interaction metaphor.

Another key finding is that as task complexity increased (from transformation tasks to modelling tasks), consensus generally became lower for both *reconfigurable* and *augmenting* PM-Objects. This could indicate that perceived affordance alone is not sufficient for participants to generate a common strategy to approach more complex tasks. Furthermore, our results from the group design workshop and user-elicited design study, showcase how participants would often discuss combining characteristics from both *augmentation* and *reconfiguration* objects. *Discrete input* is almost a universal requirement for complex tasks or modal interaction as participants would often adopt a 'clutch' metaphor for manipulation.

To summarise the takeaways around affordance mapping, we observed several agreed-upon beneficial characteristics of *reconfigurable* PM-Objects for transformation tasks, and *augmenting* PM-Objects for modelling tasks in AR. While the specific design of these interactions was not always agreed upon, we see clear PM-Objects preferences from participants in mapping affordance and designing interactions.

3.6.3 Implications of PM-Tools for AR Interaction

Our work is a first look at how PM-Tools can be used in AR and characterise the utility of *augmenting* and *reconfigurable* objects. The results from the group design workshops and elicitation study show that certain object affordances were perceived to be more appropriate for certain tasks. For example, reconfigurable PM-Objects were consistently ranked higher for transformations and augmenting PM-Objects ranked higher for certain modelling tasks. This suggests that space-multiplexed interfaces are more appropriate for complex tasks in AR environments [46], which would be more analogous to work in the real world: many bespoke tools in the 'toolbox' designed for specific tasks. Current state-of-the-art AR interaction approaches have a 'catch-all' in the form of controllers or hand gestures, and perhaps PM-Tools can facilitate more experiential interaction through an augmented or highly configurable form factor offering bespoke affordances for specialised tasks. Moreover, PM-Tools could have the potential to bridge current predominant interaction paradigms in AR, with *augmenting* PM-Objects mediating and supporting more verbose and expressive input (akin to natural gesture), and *reconfiguring* PM-Objects providing means for efficient, granular, and precise input (akin to controller interaction).

In our method section, we introduced the design space of PM-Tools, highlighting the unexplored potential of highly reconfigurable and augmenting interaction devices. The workshop results and gesture consensus set show that participants clearly appropriated different objects with different reconfiguring and augmenting properties for different tasks. As such, there are benefits to combining PM-Tools characteristics for complex spatial tasks that demand intricacy yet facilitate expression such as 3D modelling. Our conceptual examples are a first characterisation of highly reconfigurable/augmenting PM-Tools that leverage the granularity and precision afforded by reconfigurable objects and the expressiveness of augmenting objects.

This initial characterisation of PM-Tools for AR, opens up new questions for future work. Firstly, our results showed that there was no common consensus on input-output coupling for the given AR tasks with the provided PM-Objects. 34.35% of elicited gestures utilised some form of proxy interaction, either rotational, translational, or full proxy, being mostly used in transformation tasks. However, the level of I/O coupling described by participants varied considerably depending on the task, object, and user preference. 26.98% of elicited gestures involved spatial-non proxy interaction, which was mostly for modelling tasks, with participants often using the PM-Objects orientation or position to define an area for manipulation.

As with transformation, the level of I/O coupling participants described was not consistent with some participants preferring to perceive the PM-Objects as a tool acting upon a virtual object, or the PM-Objects behaving as the virtual object. Future work would need to explore I/O coupling further considering: particular AR tasks and applications; demographically diverse users; accessibility; and different types of PM-Objects. However, currently, it seems desirable to support both direct and indirect I/O coupling or provide mechanisms for users to control the level of spatial coupling between a PM-Tool and virtual objects to suit a particular task.

An approach to addressing the issue of I/O coupling is the use of physical objects in tandem with the physical environment specifically surfaces. *Spatial permanence* is an interesting advantage to PM-Tools and there are many spatial models outside of AR literature, such as the situative space model [303] and proxemic interaction [27, 138]. As such, PM-Tools and tangibles, in general, could help to promote a better understanding of space in virtual environments. For instance, tracking when an object is inside or outside the field of view of a user, tracking the distance between one or more objects and users, and identifying when objects are within the interaction space of a user. This coupled with the advantages of PM-Objects explored in this chapter, could enable interaction devices to adjust the level of I/O coupling, shape-change their affordances, and even change the mode of interaction as users move throughout a virtual space. Additionally, there are opportunities for modal interactions when an object is in mid-air or on a surface. Despite explicitly allowing participants to use surfaces when designing gestures, all of our consensus set gestures were designed for mid-air interaction. This could be due to the commonly observed 'legacy bias' in open-elicitation: the recognisable nature of the objects results in objects not being used with surfaces. It could also be due to different objects being selected for certain AR tasks so the total number of interactions for one object was limited and could all be achieved in mid-air. Participants were also not considering the gestures for prolonged use, where using a surface to avoid fatigue would be useful. Likewise, the chosen AR tasks may not be conducive to interaction with PM-Objects on a surface.

Furthermore, our work primarily focused on one user acting upon one PM-Object with one virtual object, but it could be beneficial to have physical devices work in tandem with natural gestures or with each other to offer both precise input in one hand and verbose input in the other. Asymmetrical interaction with multiple input modalities has previously been explored for tablet interaction [170], but it is unclear if and how asymmetry would work with PM-Tools and gesture, or a series of PM-Tools multiplexed together. Moreover, our results have shown how certain properties and affordances are more desirable for certain tasks, hence it could be important to consider how PM-Tools can be made modular with different tools for AR interaction constructed at the point of need. Utilising PM-Tools as virtual multi-tools is an interesting prospect and could help address the lack of 'reflection' and 'recovery' built within tangible interaction [178]. Our work only considered user actuation, but PM-Tools that support a combination of user and computer actuation could provide both rich input and output and even self-reconfigure depending on the user's task. Finally, how PM-Tools affect collaboration in AR environments is also unclear, and if different types of PM-Tools could affect the division of labour or collaboration styles. Simply stated, all of these areas need much further exploration concerning PM-Tools and general tangible interaction in AR.

To summarise the takeaways around PM-Tools for AR, our consensus set of gestures demonstrates a wealth of interactions that can be realised using PM-Tools that go beyond virtual object proxies that are so prevalent in related work. Our conceptual examples are also a first characterisation of what PM-Tools are possible by incorporating highly *reconfigurable* and *augmenting* characteristics. These sorts of devices are individually powerful when considering specific AR activities such as modelling and manipulation, but when multiplexed together they contribute to a wider AR interaction ecosystem - a 'toolbox' for Virtual Environments. We also highlight remaining challenges outside the scope of our exploration, such as I/O coupling, and speculate on potential avenues of future work such as leveraging spatial permanence of tangibles, using physical surfaces for modal interaction, and the confluence of natural input and PM-Tools.

3.6.4 Limitations and Future Work

There are some limiting factors in our work that require separate research or future studies to expand on. Firstly, our work primarily focused on a subset of 3D interactions in AR, however, there are a vast amount of interactions beyond this and PM-Tools could be applied in a similar manner. For example, *spatial permanence*, as described in the framework, can inform new interaction design beyond egocentric frames of reference in which the spatiality of an object can become intrinsically tied to interaction and output.

Further to this, our user-elicited design study was conducted with a tabletop present, however the majority of the designed gestures – especially gestures from the consensus set – did not utilise the physical surface (only 7.38% of gestures used a surface). This could be due to the brief nature of the AR tasks, only requiring around 10 to 20 seconds of physical movement. There could be instances where PM-Tools and the physical environment could be used in tandem to facilitate new interactions, for example for prolonged or more precise interactions tangible devices could be designed to utilise the stability of surfaces, but further exploration is needed. Beyond interaction design, there needs to be a wider investigation of the application domains of PM-Tools in AR. We suspect that combining mechanisms for precise and expressive input in one or more PM-Tools devices could be beneficial for activities that are simultaneously expressive and precise, such as sculpting or 3D modelling, but an in-depth exploration and evaluation is needed.

Regarding our user-elicited design study, the consensus set had a much lower representation, 193 out of 560 points (34.46%), compared to previous elicitation studies with traditional hand gestures which had 50-60% representation [308, 419]. This could suggest that user-elicitation might not translate as well to designing interaction with physical objects, in which case new methodologies are required to better support interaction design with tangibles. We also do not have insight into participants' familiarity with the study objects, so it's unclear how previous interactions with the objects and participant 'legacy bias' influenced their perceived utility for AR interaction.

Considering future work, this first exploration of PM-Tools for AR has resulted in some example interaction devices based on a gesture consensus set. To fully understand the cognitive and physical effects of these devices and evaluate their practicality, we need to test them against other forms of AR interaction using functional prototypes. Additionally, creating prototypes based on our findings would allow us to investigate challenges related to I/O coupling and spatial permanence.

3.7 Chapter Conclusion

In this chapter, we explored the application of PM-Tools to a range of AR activities, specifically 3D manipulation and modelling of a virtual object. We answered **RQ1** and the sub-research questions by developing a novel classification, vocabulary, and design space for physically-modifiable objects (PM-Objects) as a vehicle for exploring physically-modifiable tools (PM-Tools) for AR. Leveraging this framework, we describe the mapping between a set of 10 selected PM-Objects and AR interactions through a group design study and guessability study. Our results from these studies show how physical objects with reconfigurable or augmenting traits can be applied in AR for complex interactions such as virtual object manipulation, editing, and modelling. There are a number of conclusions drawn from this work:

Firstly, we found that *reconfigurable* PM-Objects had characteristics that were agreed upon for virtual object transformation tasks such as granular expansion and rotation. Likewise *augmenting* PM-Objects had characteristics that were agreed upon for modelling tasks such as organic deformation. In both types of tasks, discrete input on a PM-Object was seen as useful for employing ‘clutch’ mechanisms.

Secondly, while the choice of which PM-Objects best suited an AR task was often mostly agreed upon, the design of the specific gesture and actuation of the PM-Object was much more inconsistent. Gesture designs became less agreed upon as tasks became more complex and abstract, for example, transformation tasks compared to modelling tasks, which was especially prevalent for *reconfigurable* PM-Objects.

Thirdly, it is uncertain how I/O coupling should be designed between a PM-Object and a virtual object. Some gestures use PM-Objects as direct 1-to-1 proxies, while others treated the PM-Objects as tools to enact on virtual objects. To provide users with flexibility, it would be beneficial to offer options to virtually modify the PM-Tool’s relationship with the virtual object, such as switching the PM-Tool between a direct representation of a virtual object to a tool to manipulate the object.

Fourthly, our study designs enabled participants to incorporate their physical environment when creating gestures using the PM-Objects. Participants utilized the desk surface to support their arms during the gestures but did not actively use it as part of the designed gestures. As a result, most interactions could be executed without a surface. This makes the PM-Tools more flexible, not relying on a physical surface but could be problematic for prolonged use. To address this, we should still consider incorporating the use of physical surfaces into the design of PM-Tools.

Finally, our work showed how different properties and affordances were useful for different AR tasks. Considering this, we discussed how PM-Tools could be made **modular** and **constructed** at the point of need when working inside AR. Additionally, we also described how treating PM-Tools as multi-tools for virtual work could be an avenue to address the lack of 'reflection' and 'recovery' when purely relying on an interaction devices' perceived affordance [178].

While we presented a set of conceptual devices based on our findings, these were not exhaustive and there are still many interesting interaction concepts to be explored. In the next chapter, we delve deeper into interaction concepts such as blending natural gestures with tangible interaction, utilizing an object's 'spatial permanence' to leverage the physical environment, and incorporating multiple PM-Tools in a broader device ecosystem. Furthermore, while this chapter examined a breadth of form factors for AR interaction, Chapter 4 delves more in-depth into one specific tangible form factor — cubes. From our exploration of existing physical objects for AR tool-making, cubic form factors appeared to be particularly prevalent and, on inspection of related work [235, 339], have unique benefits for interaction design as well as opportunities for scalability and modularity which we describe further in Chapter 4 Section 4.1. Finally, we also aim to address some of the non-trivial technical challenges such as fabricating and prototyping an instrumented modifiable tool, detecting user input, and tracking the tool in 3D space.

Chapter 4

Tool-Modifying

The previous chapter showed how particular affordances of physically-modifiable objects could be appropriated for activities in AR, mediated by a tangible interaction device. We found that gesture design of a particular interaction was not always agreed upon and we discussed other potential interaction concepts that require more exploration. More specifically, these were combining natural gestures and physical tools, leveraging the physical environment (surfaces), and multiplexing physical tools together in AR. In this chapter we focus on exploring this design space further with one specific tangible form factor — cubes — to mediate interaction in AR to address the second research question (**RQ2**): *How should physically and virtually modifiable tools be created and operated for Augmented Reality?* There are several sub-questions to **RQ2** categorised into two groups based on design and technical challenges:

- **RQ2.1:** *Design challenges and opportunities of modifiable-tools in AR*
  - How can a physical tool be virtually modified in AR?
  - What are the design opportunities for multiple physical tools in AR?
  - What are the design opportunities for combining a physical tool with user gesture?
- **RQ2.2:** *Technical challenges of modifiable-tools in AR*
  - How can we fabricate a physically-modifiable AR tool?
  - How should the tool be instrumented in order to detect user gestures?
  - How do we track the AR tool in 3D space?

In this chapter we focus on designing physically-modifiable AR tools (PM-Tools) using one specific form factor — a **cube**. Having explored a number of PM-Objects in the previous chapter, cubic form factors were particularly prevalent in the representative set of *reconfigurable* PM-Objects that were used in the two studies in the prior Chapter. In addition, the benefits of cubes as tangible interaction devices have previously been explored [235, 339]. Particularly the work of Lefeuvre et al. [235] categorises the distinct affordances and properties of cubes by "manipulation, placement in space, arrangement, multi-functionality, randomness, togetherness, physical qualities, containers, and pedestal for output". These benefits have also been explored in a number of different VR and AR environments [17, 117].

To this end, the goal of this chapter is to instill the principles of physical and virtual modification into a form factor that is conceptually, technically, and physically as simple as a cube. In order to address **RQ2** and the sub-questions, we take inspiration from previous work on tangible prototyping kits and present a first exploration of the emergent design space of combining cubes and their beneficial affordances with AR and surface-based gestures. We then detail the design of the cube tool, the hardware implementation and fabrication process, and the components of the toolkit. Following this, we evaluate the gesture detection capabilities of the toolkit through three demonstrative applications, displaying the range of surface-based gestures supported and highlighting the generality of the toolkit. Next, we extend the toolkit to incorporate spatial tracking of the cubes in 3D through the principle of tracking the user's hands (something which is available in most AR/VR HMDs). Finally, we reflect on and discuss the features and implications of the toolkit. Part of the work in this chapter was originally published in the Sixteenth International Conference on Tangible, Embedded, and Embodied Interaction (**TEI'22**) [310].

4.1 Augmented Reality & Cubes

Consider the design space of combining cubic tools with AR. Cubes are a promising TUI device, with a history of use as controllers, tokens, and interactive modules. The design space of cubic tools in AR is guided by the '*Bricks, Blocks, Boxes, Cubes, and Dice*' taxonomy of design properties of cubes which were developed by surveying research, books, and products on cubic interactive tangibles. More specifically, we focus on a subset of themes from the taxonomy for exploring the design and interaction opportunities in AR: *manipulation as Input*, *placement in space as input*, *Arrangement*, *Containers*, and *Pedestal for Output*. In this section, we describe cube affordance in AR in light of this taxonomy and introduce four different *interaction metaphors* to apply to cubic tools in AR which include: *Cubes as a Spatial Proxy*, *Cubes as a Container*, and *Cubes as a Manipulation Handle*. Each interaction metaphor describes how a cube can be virtually modified to provide different functions in an AR environment.

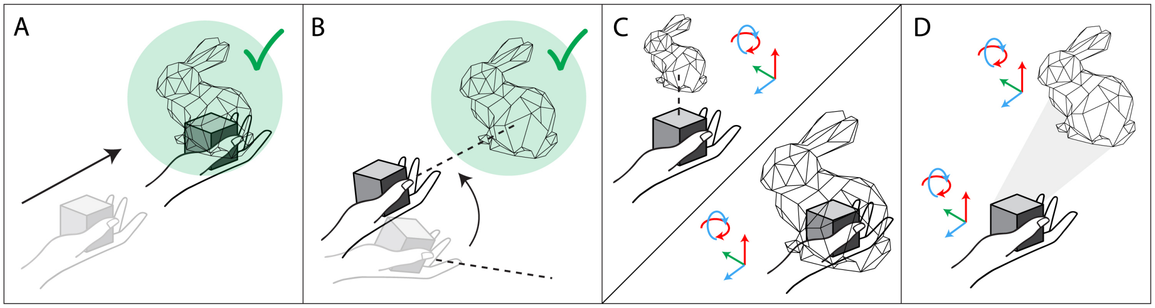

Figure 4.1: **Cubic tools as a proxy**: **A)** Direct selection of a virtual object is performed by physically intersecting a proxy cube with the desired virtual object. **B)** Remote selection is performed via raycast, intersecting the ray with the object, and confirming the selection. **C)** Exact spatial manipulation is performed after a direct selection in which the cube directly manipulates a virtual object's translation and rotation either from an attached point on the virtual object or through WIM. **D)** Relational manipulation is performed after remote selection in which the cube indirectly manipulates the virtual object's translation and rotation while maintaining the relation from the point of selection.

4.1.1 Cubes as a Proxy

Using cubes as spatial proxies for virtual objects is the primary mechanism for performing 3D manipulations on virtual objects, specifically providing 6 degrees of freedom (6DoF) for translation and rotation. This interaction metaphor operationalises '*manipulation as input*' and '*pedestal for output*' from Lefeuvre et al.'s cubic tangibles taxonomy [235]. As mentioned in Chapter 2, Tangible Objects in AR have widely been used as proxies for virtual objects due to facilitating seamless continuity between displaying and interacting with virtual content, as well as leveraging passive haptics for improved motor-control and immersion [46]. Hence, this is something that should be retained for a cubic tool in AR due to the strong affordance for manipulation. A cube as a spatial proxy for a virtual object can be designed in a number of different ways but, as described by 3D UI fundamentals, there must be mechanisms for selection and manipulation:

3D Selection: Before transforming a virtual object, it must be selected from a scene or set of virtual objects [221]. Different approaches and challenges exist for cube-based selection. One approach is to intersect the physical cube with a virtual object, using collision as a means of direct selection. Remote selection, i.e. virtual objects that are at a distance, requires virtual modification of the cube to support pointing via raycast or combining it with other inputs such as hand gestures, voice, or gaze. Confirming a selection usually depends on a discrete input mechanism (usually a button in controllers) which can be supported by surface-based gestures on a cube face. Issues such as target disambiguation and occlusion in 3D remote selection can be addressed by combining raycast functions applied in controllers, such as bubble cursors [28, 141], with control-display gain adjustment via surface-based gestures [68, 312, 317].

3D Manipulation: Cubes can serve as direct proxies for virtual objects providing a 1-to-1 spatial mapping. An ”exact” mapping is achieved by overlapping a cube and a virtual object, either by projecting the virtual object as a miniature onto the cube (WIM [358]), or by attaching the cube to a point on the virtual object. A ”relational” mapping is where a cube and virtual object's positions are relative, allowing for 3D manipulation from a distance or at a larger scale. Generally for rotation, a cube's ability to be manipulated in 360° using only one hand [339] is potentially beneficial for supporting complex manipulations with virtual objects such as mixed transformations. This is something that is not afforded in standard AR/VR controllers and we refer to this as a cube's ability to be *regrasped*.

Furthermore, for precise manipulations, a cube's shape clearly represents the three spatial axes (x,y,z) [221, 235] leading to a more intuitive understanding of DoF separation [155]. Precision can be further improved by using dynamic gain control through a combination of two cubes, where one cube controls the other, which we expand on in Chapter 5. Another way to increase manipulation precision is to combine the cube with a physical surface to function like a mouse allowing a virtual object to be dragged around. Realizing these types of manipulations requires both cube multiplexing and discrete input mechanisms, such as surface gestures, for spatial clutching.

In summary, using a tangible object as a spatial proxy has been proven to be an intuitive interaction in Tangible AR and has been explored thoroughly in related work. However, cubic tools specifically offer a wide range of capabilities and potential interactions for 3D manipulation by utilizing their form factor, physical surfaces, multiple cubes, and surface gestures on the faces. In Section 4.6, we showcase two examples of this interaction and in Chapter 5, we provide a comprehensive overview of a 3D manipulation technique in AR.

4.1.2 Cubes as a Container

The concept of a Tangible ‘container’ is a central aspect of early TUI research in the 1990s and early 2000s [189]. For example, Ulmer et al.'s work on MediaBlocks [388] proposed the idea of physical containers that do not actually store digital information but instead provide a reference or dynamic association to the intangible content such

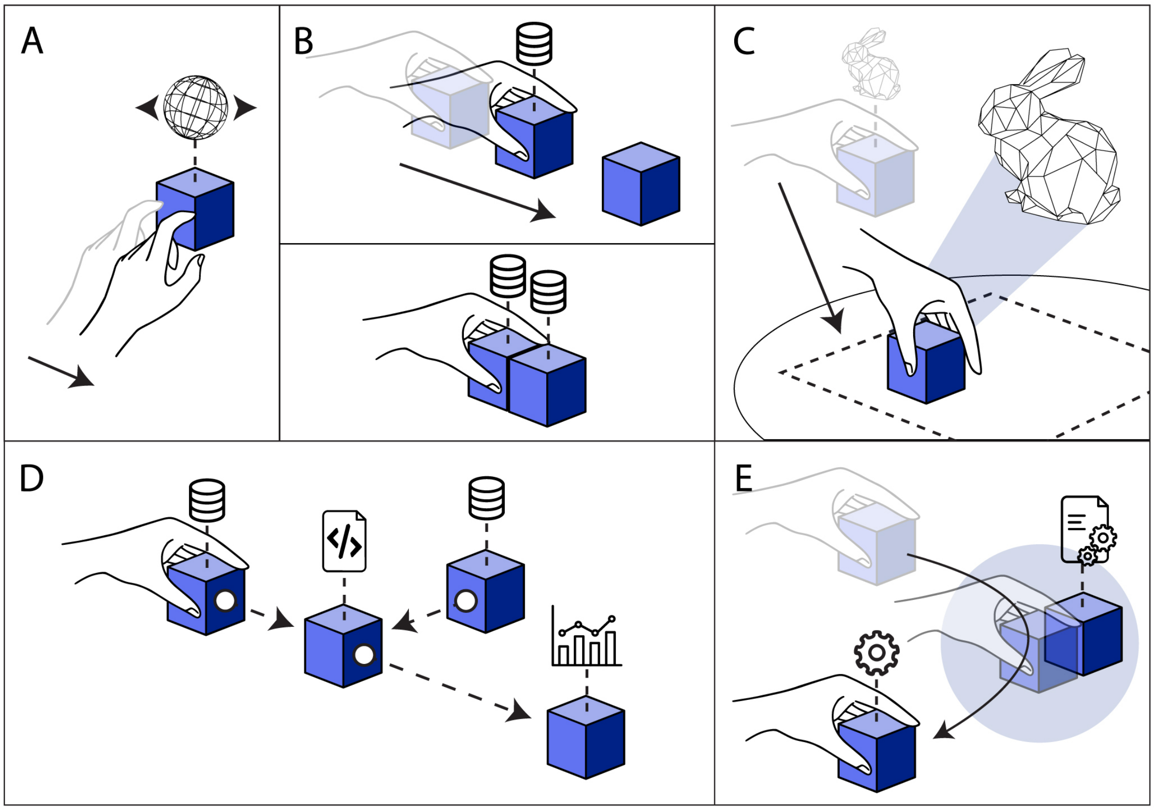

Figure 4.2: **Cubic tool as a container**: **A)** A user can cycle through the items stored within a container by utilizing the 3D thumbnail displayed above it. **B)** Users can copy an item from one container to another by spatially arranging the cubic tools next to each other. **C)** To instantiate an item within the virtual environment, a user can place the container directly into the designated staging area. **D)** The containers can be spatially arranged in an application pipeline, where two containers hold datasets that are processed by a middle container, resulting in a visualisation that is stored and displayed in a final container. E) An example of a container used as a 'meta instrument' [31] in which a blank cubic tool has its function configured by being spatially arranged next to the container.

as data and media. Typically containers in TUI research are utilised as mediators between media sources (images, video, audio, etc.) and display devices (projectors, printers, etc.). More recent work, as reviewed in Lefeuvre et al.'s taxonomy [235], has proposed cubes to also be 'pedestals for output' (display devices) and media sources such as in the work of Jordà et al. [199]. Combining the increased display capabilities and 3D tracking of AR HMDs with cubic tools as 'containers', yields new possibilities for interacting with media and information, providing users a means to physically arrange their virtual workspace of intangible content. As a result, this interaction metaphor leverages cubes as '*containers*', '*pedestals for output*', '*placement in space*', and '*spatial arrangement*' from Lefeuvre et al.'s taxonomy [235].

In AR, media can be displayed in a 3D environment along with virtual models, scenes, and visualizations, much like in desktop interfaces where we manipulate documents, images, and applications. A cubic tool functioning as a container in AR hypothetically has boundless storage for heterogeneous media and information. Containers should also facilitate interactions we are familiar with from desktop interfaces such as moving, copying, and deleting information [349] from a container.

There are a number of approaches to designing interaction with cubic containers in AR. For communicating what a container is storing, we can utilise a cube as a '*pedestal for output*' and use 3D thumbnails directly on or above the containers showcasing WIM [358] style models of virtual objects, visualisations, images, or documents. Copying and moving information from one container to another can be as simple as using the '*spatial arrangement*' of two containers next to each other to move content from one to the other. Gestures, such as motion or surface-based, can also be utilized for functions such as cycling through container content, swiping content between containers, or clearing a container by shaking it.

'*Placement in space*' can also be appropriated by containers in AR to seamlessly add content to a virtual environment. For example providing designated areas in the physical environment where once a cube is placed, content held within the container is loaded into the virtual environment. This allows a seamless connection between WIM interaction [349] with content in a container (e.g. with the 3D thumbnails), and direct interaction with virtual content that is loaded into the scene from the container. We can also take inspiration from related work on cross-device interaction [59] and proxemic interaction [27] when designing containers for transferring content between display areas (or staging areas in AR) and individual containers.

In addition to storing media, cubic containers in AR can also hold entire virtual scenes, AR applications, and instructions for configuring other cubic tools. Due to the boundless nature of containers, any number of virtual objects can be grouped together and stored as a virtual scene, allowing for the saving and storage of an entire virtual workspace for future use. Cubic containers can hold running applications and be arranged to manipulate application parameters, similar to how tokens in the

ReacTable are used for music production [199] and how tangible programming blocks are assembled to create program functions [176, 184]. In this way, physical data pipelines can be set up using the containers to manipulate data in AR. For example containers with different datasets at the start of a pipeline act as input to other containers, containers in the middle can filter or transform data, and containers at the end display the final 3D visualisation.

Finally, a container in AR can store instructions or program scripts to modify other cubic tools in the device ecosystem, allowing a user to configure a blank cubic tool by placing it next to the relevant container. This is inspired by Beaudouin-Lafon's 'Instrumental Interaction' specifically relating to '*Meta-Instruments*': the idea that tools themselves can be objects of interest' and 'meta-instruments being used to organise instruments in a workspace' [31]. Cubic containers can operate as 'meta-instruments' virtually modifying the roles of other cubes in the AR environment. This is further expanded on in the next Section.

In summary, the use of containers in TUI has been extensively studied and offers ample possibilities in AR [189, 388]. By combining the attributes of cubes and AR HMDs and incorporating user gestures, the concept of *cubes as containers* can be further explored for various AR applications. We will delve deeper into these possibilities in subsequent sections.

4.1.3 Cubes as a Handle

As LaViola et al. explains [221], 3D interaction includes system controls to interact with virtual objects' abstract components, such as color and texture, or input parameters, such as text entry. In VR/AR, system controls are typically achieved through a combination of pointing (via hand controllers or gestures) and menus/widgets. In Tangible AR, physical controls have been explored for system control, such as 'Opportunistic Interaction'[156, 157, 159], which uses the physical environment. System control has even been explored for cubic form factors such as the work of Van Laerhoven et al. [393]. The '*cube as a handle*' interaction metaphor refers to using a cubic tool for system controls in AR and can be designed in a number of ways. We propose three types of cube handles: *spatial*, *gestural*, and *hybrid*.

Spatial handles control AR parameters by moving the cube in space, either absolutely or relative to other cubic tools. This utilizes a cube's affordance for *arrangement* and *placement in space* [235]. For example, a cube on a surface can act as a slider to increase or decrease a virtual object's scale or can be used to define the boundaries of a virtual area.

Gestural handles control parameters through motion gestures (such as shakes) or surface-based gestures (such as taps and swipes) using a cubic tool. This utilizes

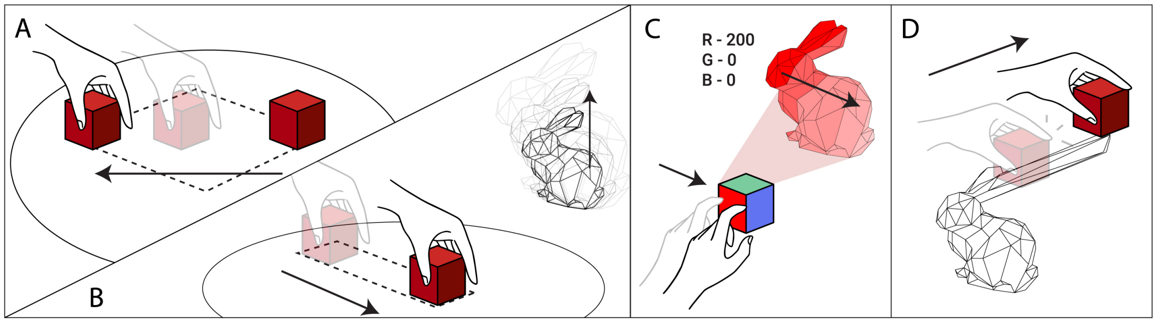

Figure 4.3: **Cubic tool as a handle**: **A)** A cubic tool used as a spatial handle to define a virtual workspace area, and **B)** as a slider to scale a virtual object. **C)** A cubic tool used as a gestural handle to change the colour of a virtual object, with 3 different faces acting as sliders for red, green, and blue values. **D)** A cubic tool used as a hybrid handle in which the cube intersects a point on the virtual object's mesh, a tap gesture selects the point, and the point is extruded to a desired position.

the *multifunctionality* property from the taxonomy, such as providing discrete touch input on certain faces. For example, using the cube to control a virtual object's color, with three cube faces acting as touch-based sliders for red, green, and blue values.

Hybrid handles use a combination of spatial and gestural input to control parameters. For instance, a virtual object's mesh can be stretched and extruded by aligning the cube over the mesh, tapping a face to select the part of the mesh, and moving the cube and mesh in the desired direction.

For all handle types, their function can be communicated in AR using the cube's affordance as a *pedestal for output*. This can be achieved by projecting a tool thumbnail or description over the physical cube. Additionally, handles can also be combined with other cube interaction metaphors, proxies and containers, to enable new interactions or improve performance:

Combining Handles and Proxies: Combining two cubes for proxy interaction can increase user expression and precision. One cube can act as a handle while the other serves as a direct proxy for the virtual object. The work of Hinckley et al. [170] demonstrates the benefits of hand asymmetry when performing pen-based interactions on a tablet, with one hand providing fine grain control while the other performs more coarse movements. The same principle can be applied by combining cubic handles and proxies in a bimanual technique. One cube can provide direct manipulation of the virtual object, while the other allows for control over the coupling between the cubic proxy and the virtual object, such as adjusting the c/d gain or toggling DoF constraints. A more detailed example of this principle is implemented in Chapter 5.

Combining Handles and Containers: Combining a handle and container cube enhances explicit control over container functions, such as cycling items, deleting content, and duplicating content, essentially enabling ad-hoc construction of virtual workspaces. For example, a cubic container placed on a physical surface might hold a virtual scene, and a handle placed next to it can be used to configure the boundaries of that virtual scene when it is spawned into the virtual environment. For more complex containers, like those containing applications, handles can be used to configure input/output faces, load/unload applications, or manipulate input variables in a container pipeline.

Using Handles as 'Meta-Instruments': Handles can also be used to modify other cubic tools such as changing one cube from a proxy, to a container or to a handle. As discussed, this can be accomplished by using containers that store different cubic tool functions and by placing a blank tool near the container. Alternatively, a handle can be defined by a virtual environment area where once placed, a cubic tool will be automatically modified to a different function. Additionally, by utilizing a cube's *multifunctional* affordance, different faces on a handle can be mapped to different tool functions, and placing another cube on one of these faces will modify its function, similar to a configuration port.

In summary, cubic tools can serve '*as handles*' to provide mechanisms for system control. Handles can also be used in tandem with other interaction metaphors to enable new types of interaction or increase performance by utilising the benefits of bimanual interaction. As with containers, handles can also be used as 'meta-instruments' [31] to configure other cubic tools being used in the AR environment.

4.1.4 Summary and Technical Challenges

To summarise, we have partially addressed **RQ2.1**. There are a number of different *interaction metaphors* that can be applied to cubic tools in AR, inspired by the wealth of literature on TUI [338] and the unique affordances of the cube form factor [235, 339]. While other form factors certainly have the capability to support these kinds of interaction metaphors, cubes have the unique capacity to support all of them simultaneously when multiplexed across several identical devices. Moreover, a cube's stability when positioned and arranged on physical surfaces makes it an ideal starting form factor to explore physical surface-supported activities in AR. We also argue that a cube can be used as a starting point to explore the virtual modification of tools and over time through empirical study, more specialised form factors may arise that are better suited to certain functions.

Reflecting on the design space of cubes in AR, there are a number of sensing and tracking challenges that need to be addressed in order to realise the proxy, container, and handle interaction metaphors. Firstly, it is crucial that the cubes can be tracked in 3D space without relying on computer vision. The interaction metaphors described rely on the assumption that an AR environment will have an understanding of where the cubic tools are positioned, and the cubic tools themselves will be able to sense one another. Tracking the cubic tools using computer vision is problematic due to the user's hands or other cubes occluding one another and commercial AR HMDs perform computer vision tracking only from the user's point of view, making it problematic when performing indirect interactions with the cubes (i.e. when a virtual object is 'over there'). To address these challenges, we detail the design of a prototype cubic tool in Section 4.3 and describe how 3D object tracking can be achieved without entirely relying on computer vision in Section 4.5.

Another technical challenge involves detecting touch gestures on the cubic tools themselves, as several interaction metaphors require this input. However, before delving into the technical challenges, it is necessary to outline the potential interactions that can be achieved by combining a cubic tool with touch-based surface gestures, which will be the focus of the next section.

4.2 Cubes & Surface Gestures

After describing the interaction metaphors that can be applied to a cubic tool in AR, now we look at the potential interactions that a user can perform on a cube itself. The design space is divided into surface-based input, combining gestures, interactions beyond gesture, and output space.

4.2.1 Surface-based Input

The design space of single and multi-touch surface gestures is based on cubic tangible and surface-based computing literature [235, 419]. The surface interactions described in Figure 4.4, while not an exhaustive list, demonstrate the range of touch gestures that can be performed on a simple cube: single touch gestures such as taps, swipes, and path gestures, and multi-touch gestures such as multi-finger taps and pinch gestures. Due to the stable form factor of cubes and the inherent graspability, these gestures can be performed in hand or while the cube is at rest. By combining cube affordance and simple surface-based gestures, common TUI metaphors can be mimicked: swipes to represent a slider, taps as button input, path traces as directional input or as a dial, and pinch gestures to replicate common touchpad input.

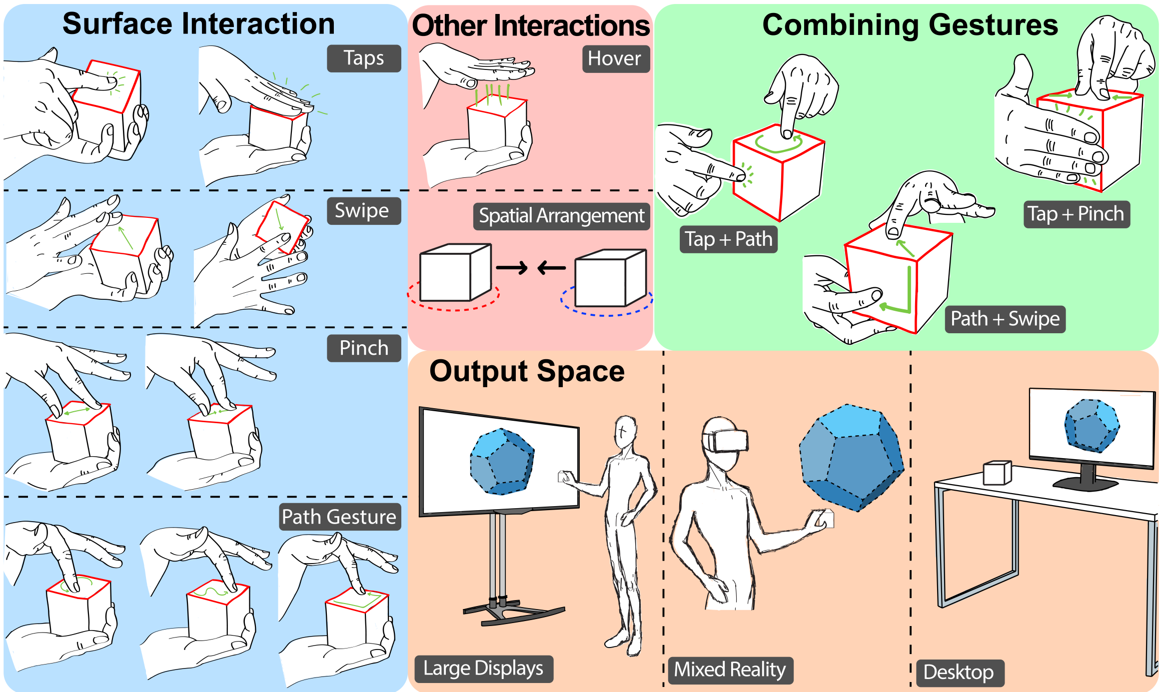

Figure 4.4: A visualisation of the interaction and output space combining cubes and surface-based gestures.

4.2.2 Combining Gestures for Input

Cubes are multiplexers of interaction by virtue of their form factor. Each face, while appearing identical to one another, can be leveraged as a separate area for input and even configured on the fly depending on the context of interaction. In the case of surface-based gestures, a cube can support simultaneous and heterogeneous surface-based gestures on any face. Further to this, a cubic tangible can support modal or state-based interactions depending on whether the cube is on a surface, in hand, or actively touched on a particular area.

4.2.3 Other Interactions

There are multitudes of other interactions that are significant for the design space of cubic tangible interaction. Based on the work of Lefeuvre et al. [235], cube affordance alone has interesting properties such as manipulability, spatial arrangement, and multi-functionality which we discussed in the previous section. However, more interactions are capable when you instrument an object using a particular sensing approach. For example, capacitive sensing enables the detection of proxemic interactions, such as non-surface-based gestures, or even other capacitive devices.

4.2.4 Output Space

Depending on the interaction context, different advantages of cube affordance can be leveraged in addition to surface-based interactions. As described previously, a cube's dimensionality, manipulability, and ability to be a pedestal for output make them ideal candidates as interaction devices in AR. While we just focus on AR HMD interaction, combing the benefits of cube affordance with gesture detection can produce a wide variety of unique interactions across a number of display mediums.

Moreover, the inherent tangibility of cubes makes them an ideal interaction device to support multi-user collaboration, having the ability to be freely passed from one user to another, be placed in the physical environment, or even act as a mediator when users might not have the same display capabilities (for example one user in VR, one in HMD AR, and another using mobile AR).

As mentioned, cubes are also suited in output spaces that leverage physical surfaces, such as AR, desktops, or tabletops. The stability and space afforded by surfaces not only allow users to easily arrange and configure cubes but also combine surface gestures simultaneously on different faces. In this case, a cubic tool can transition from a manipulable object in hand, to an in-situ and fixed controller.

4.2.5 Summary and Technical Challenges

We have addressed the final part of **RQ2.1** by highlighting the number of interaction possibilities that exist when combining cubic tools, surface-based gestures, and AR display capabilities (see Figure 4.4). For **RQ2.2** however, sensing a wide array of surface-based gestures without adorning the user or using external tracking is complex to fabricate and often lacks scalability. Instrumenting tangibles with capacitive sensing is a popular approach in related work [335], however, gesture recognition using capacitive sensing is also non-trivial and requires different sensing configurations depending on the surface gesture being detected. To address this we developed a toolkit that provides a method for fabricating modular capacitive cubic tools that detect distinct on-surface gestures which we describe in the next section.

4.3 Toolkit Design

To address key challenges that exist in fabricating interactive cubic tools and to explore the questions in **RQ2.2**, we developed a toolkit which we refer to as *TangibleTouch* [310] (Figure 4.5). The goal of this toolkit is to provide a rapid fabrication method that uses simple single-extrusion 3D printing to support the prototyping of modular and physically-modifiable cubic objects. The cubes have interchangeable capacitive faces with different sensor configurations designed for

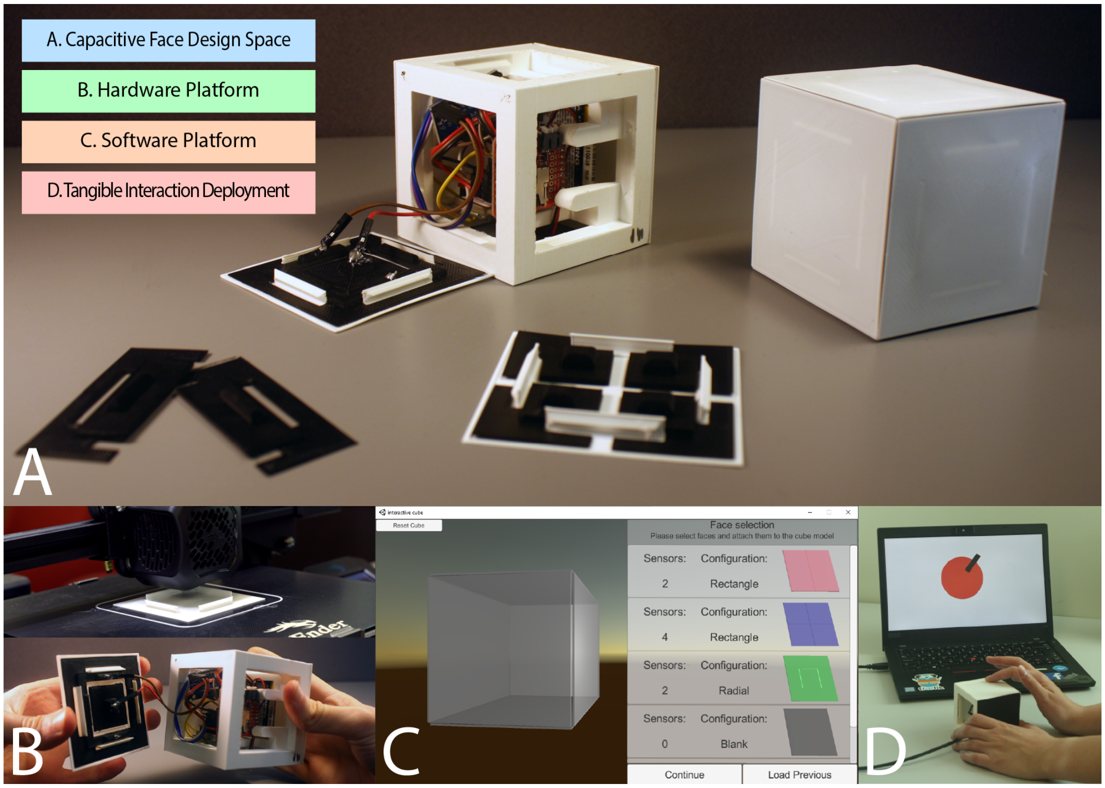

Figure 4.5: The TangibleTouch toolkit composed of: **A)** Capacitive face design space, **B)** modular hardware platform, **C)** software platform to train surface gestures, and **D)** deployable interactions to any Unity application.

surface-based gesture recognition. Additionally, the toolkit aims to provide a software platform for digitally configuring the faces of a cube, training a particular face to detect one or more surface gestures using machine learning, and a means of deploying those interactions in a variety of interaction contexts and applications. The toolkit consists of 3 parts:

1. A face design space for divvying up the surface area of a cube to design for single and multi-touch gesture recognition using capacitive sensing.

2. An extendable hardware platform using conventional 3D printing that allows for interchangeable faces with different capacitive configurations.

3. A software platform that provides a user interface to add and configure interactive faces to a cube, record data of surface gestures and train a machine-learning model for gesture detection. These can then be deployed as interactions to a variety of different Unity-based applications.

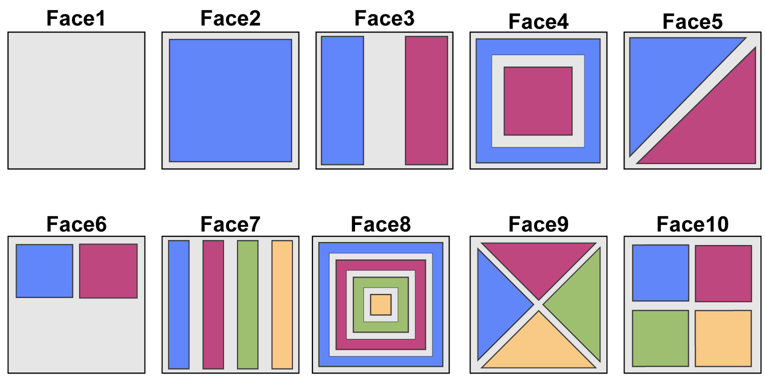

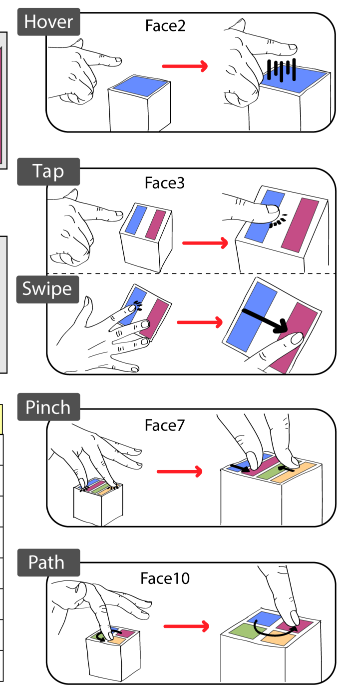

		Face1	Face2	Face3	Face4	Face5	Face6	Face7	Face8	Face9	Face10
Attributes	Sensor Number	0	1	2	2	2	2	4	4	4	4
	Configuration	Blank	Square	Rectangle	Radial	Cross	Square	Rectangle	Radial	Cross	Square
	Face Coverage	0%	100%	66%	66%	66%	66%	80%	80%	80%	80%
Interaction	Tap	X	√	√	√	√	√	√	√	√	√
	Pinch	X	X	X	X	X	X	√	X	X	X
	Swipe	X	X	√	√	√	√	√	√	√	√
	Path	X	X	X	X	X	X	√	X	√	√
	Hover	X	√	√	√	√	√	√	√	√	√

Figure 4.6: Capacitive face design space showing the configuration of sensors and the touch interactions afforded.

4.3.1 Face Design

To guide the design of the capacitive faces, we started with five general interactions that we wanted to detect using the lowest number of touch-sensitive areas: tap, swipe, pinch, path, and hover (see Figure 4.4). Figure 4.6 shows a number of different sensor configurations, varied by the number of interactive areas and their placement on a given face, followed by which surface gestures these can support. The concept for each face design follows a principle of low complexity in terms of the number of discrete touch areas. Instead, we rely on the multiplexed nature of a cube, with dedicated faces for particular interactions. While the toolkit can support more complex face designs, we focus on 10 simple face designs with 4 different sensor quantities (0, 1, 2, and 4), and 4 different face configurations: square, radial, rectangle, and cross. Figure 4.6 demonstrates how these variables determine if a surface gesture can be supported. For example, face2 supports a tap but not a swipe as opposed to face3 which supports a tap or swipe but requires double the amount of touch areas. For detecting a swipe gesture in practice, face3's sensor0 is triggered at the start of a swipe gesture and sensor1 at the end and vice versa for a different swipe direction. Generally, if more complex gestures are to be detected the number of touch areas on a single face increases, or if a face needs to support more than one surface gesture.

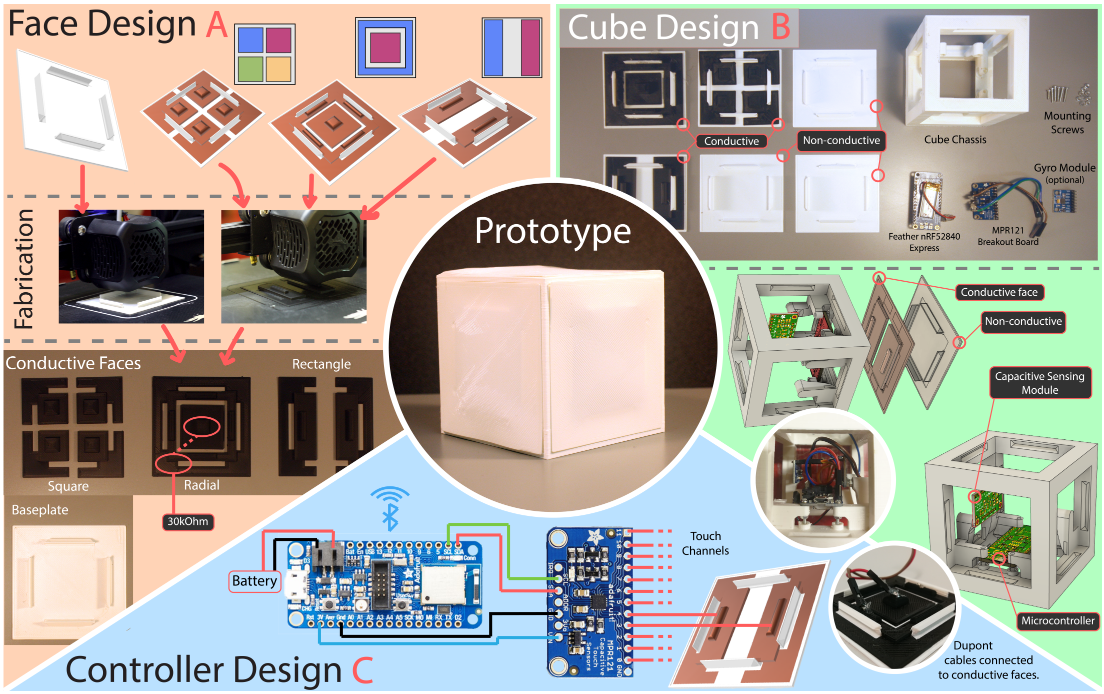

Figure 4.7: The toolkit's hardware components and design. **A)** The face design and fabrication, **B)** the cube components and modular design, and **C)** the controller design, and connection to the conductive faces.

Another example is that both face3 and face5 support swipes, but the directionality of the swipe relative to the rest of the cube would be different, i.e. face3 cannot support 'top-to-bottom' swipes whereas face5 can.

4.3.2 Hardware Platform and Fabrication

The hardware platform for TangibleTouch (see Figure 4.7) consists of three main components all of which can be fabricated using a conventional, single-extrusion 3D printer: i) non-conductive face bases, ii) conductive face components, and iii) cube chassis. In the face design, we explicitly chose to design the conductive and non-conductive parts to be printed separately to make fabrication more viable for single-extrusion printers, which are generally more accessible and commercially available, as opposed to dual-extrusion printers. We also use a capacitive sensor board, Arduino microcontroller, and lithium battery to instrument the cube.

Non-conductive face bases and cube chassis can be printed using generic PLA or ABS. The conductive components can be printed using either conductive PLA or ABS with a conductance of at least 4.6*102 Ohms/cm. For the tangible cube prototypes

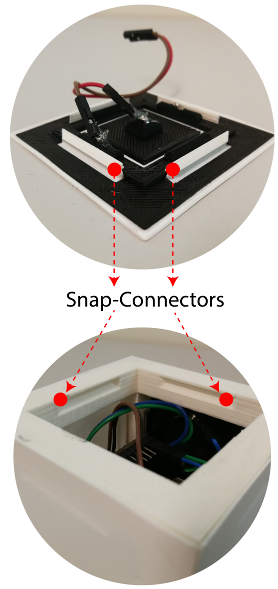

Figure 4.8: The snap connectors to mount capacitive faces.

we developed using the toolkit, the non-conductive parts were printed using white Filamentive PLA, and the conductive parts were printed using U3 conductive ABS with a conductance of 4.64*102 Ohms/cm. We recommend ABS for the conductive components as acetone can be used as a means of adhering the conductive pieces to the non-conductive base plate giving reliable adhesion with less impact on the surface capacitance. Conductive PLA can be used, and can often provide better conductivity, but an adhesive agent is needed to mount conductive components to the non-conductive face base plate, which may affect the surface capacitance. All 3D printed parts were printed on an Ender3 V2 at 50mm/s and a layer height of 0.16mm. The cube chassis took 8 hours to print and each face base plate took 30 minutes (11 hours total for non-conductive parts). Conductive part print times can vary depending on the surface coverage, from 30 minutes to 1 hour.

Dupont cables were used to connect the conductive parts to the capacitive sensor board, by heating a male connector using a soldering iron and inserting it into the mounting points. The prototype conductive components were measured at around 30kohm resistance across the conductive surface, and 10kohm from the conductive surface to the connecting cable. We also tested ProtoPasta conductive PLA mounted to the base plate using hot glue, which measured at 6kohm resistance across the conductive surface and 4kohm from the surface to the connecting Dupont cable.

Once conductive components are mounted to the face base plate and the cables have been connected to the conductive mounting points, the face can be simply attached to the cube chassis using 'snap-fit' connectors and the cables routed to the 12-channel capacitive sensor board (see Figure 4.8). In this case, a single cube can have a maximum of 12 discrete touch areas and using the modular faces can be distributed in any manner across the cube. For example, 2 faces with 4 touch areas, 2 faces with 2, and 2 with 0 or 6 faces with 2 touch areas.

We used an Adafruit MPR121 12-Key Capacitive Touch Sensor breakout board for detecting capacitance, connected to an Adafruit Feather nRF52840 Express microcontroller powered by a small 3.7v 110mah lithium polymer battery, all of which can be mounted inside the cube chassis using M2 screws and nuts. The Adafruit Feather has low-energy Bluetooth capabilities for transmitting data to other devices, which using a 110mah battery, can run for 11 hours on a single charge.

We experimented with an additional gyroscope module that also has space for mounting within the cube chassis. The purpose of this is to make configuring the cube using the software interface easier, as a designer can determine which faces map to the digital representation by simply tilting the cube. However, this is optional and is not necessary for configuring the cube.

4.3.3 Software Platform

The TangibleTouch software platform is used to digitally configure the cube with the appropriate interactive faces, train a particular surface gesture for a given face using machine learning, and deploy the trained model of a surface gesture to an application (see Figure 4.9). The software platform consists of three modules: i) An interface library, ii) a gesture-training library, and iii) a data processing and hardware library.

4.3.3.1 TT Interface and Configuration

The Unity-based interface can be loaded as a scene in a developer's application to then configure a cube, train gestures, and deploy interactions packaged as Unity events, which applications can subscribe to. On loading the interface, a designer scans for Bluetooth devices or selects the cube directly if connected via UART. Once the cube is found and selected, the designer is taken to the configuration screen. Here, a designer can see the virtual representation of the cube device with 6 blank faces and a side panel with faces of varying sensor configurations. If a gyroscope module is connected to the cube, then the virtual cube will mimic the rotation of the physical cube to decipher the face positions, otherwise, a developer can manipulate the cube using the mouse. Developers can easily add additional face configurations by invoking the face class and providing a '.obj' file.

To first configure the cube, a developer needs to select each face on the virtual model and assign a face configuration. Once all face configurations have been assigned the designer then cycles through each interactive touch area on the virtual model and touches the corresponding capacitive areas on the physical cube to calibrate the

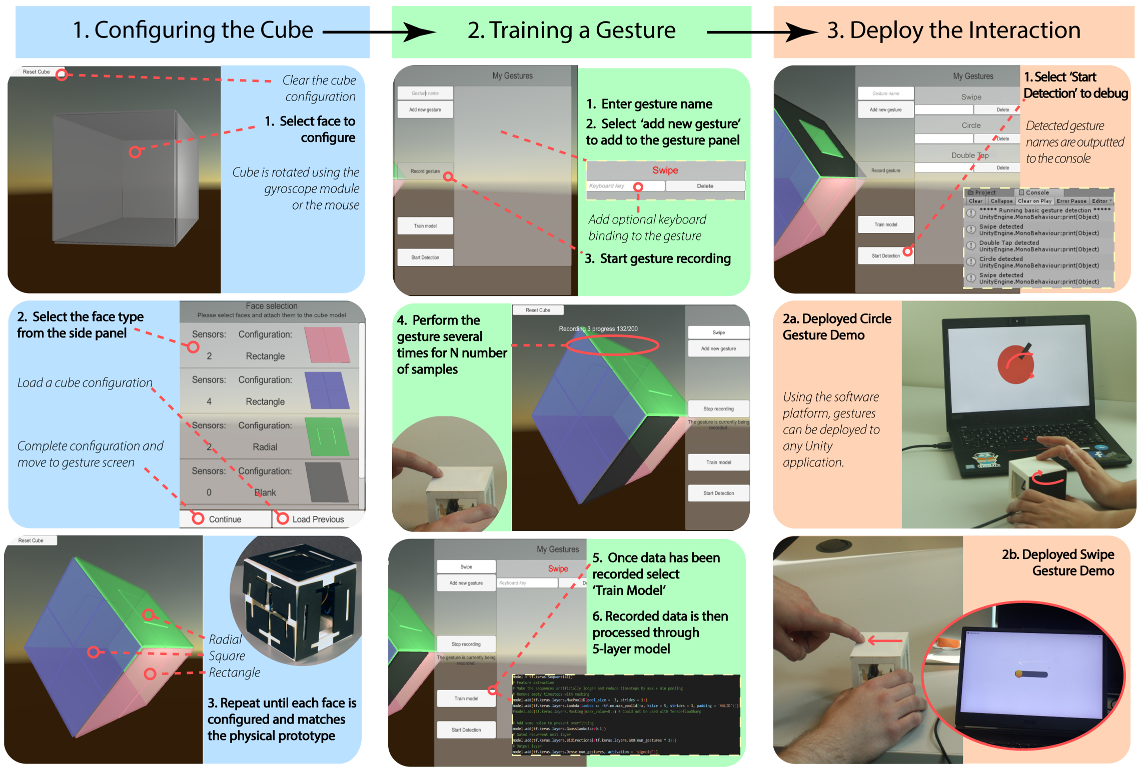

Figure 4.9: The software platform to configure, design, train, and test surface gestures including some examples of deployable interactions.

sensor channels on the capacitive sensor to the face configurations. If a previous configuration has been made, then a designer can load this into Unity from file by selecting the 'Load Previous' button. A cube configuration is cleared by clicking the 'Reset Cube' button. Once a designer is happy with their cube configuration, they can move to the gesture-training screen by clicking 'Continue'.

A designer adds a new gesture to the cube by uniquely naming the gesture and selecting 'Add New Gesture'. An optional keyboard binding can be added for that gesture that will be triggered on gesture detection. The designer then selects the newly added gesture and selects 'Record Gesture'. Now the designer performs the desired gesture 20 times, each with a sample size of 200 over a 3-second window. The number of gesture samples, the sample size, and the duration can all be altered, but these were the most optimal settings considering the time to set up and accuracy. Once any number of new gestures are recorded, the model can be retrained to include the newly added gestures by selecting 'Train Model'. 'Start Detection' then deploys the trained model, firing Unity events or triggering keyboard input depending on whether any gestures are detected. Developers can have an external class subscribe to the events fired by the gesture detection class, with each event containing a unique gesture name. Figure 4.9 shows an example of recording data for a swipe gesture and then deploying this interaction to a simple slider.

4.3.3.2 TT Data Processing and Gesture Training

For the capacitive sensor board, the default firmware settings were sufficient for the most part, but after testing it was found that the non-conductive base plate covering the conductive areas causes the baseline signal to not adjust quickly enough when a sensor is touched, which affected the performance of gesture detection. Setting the filter delay register (MPR121FDLF) to the maximum value of 255 greatly improved the touch sensitivity and baseline stability.

Gesture detection was implemented using TensorFlow [1], a recurrent neural network, incorporated into the Unity environment using a standalone Python program. Once the model is ready to be trained, recorded data is loaded from a CSV file, and any non-configured sensors are zeroed. To account for a user holding the cube or the cube resting on a surface while performing a gesture, any sensor channels that are triggered for more than 40% of a gesture sample are disregarded and zeroed. The model itself consists of 5 layers.

First, the data goes through feature extraction, consisting of maximum and minimum pool layers. These are typically used to down-sample and extract features from images by partitioning them into a set of non-overlapping rectangles and, for each such sub-region, outputs the maximum/average or minimum. In this case, max-pooling is used to increase the sequence length to smooth out any anomalous samples. As max-pooling decreases the number of lows, i.e. timesteps where the sensors are not touched, between touches, min-pooling had to be performed to increase these gaps between touches. Setting the strides to 2 also down-sampled the data from 200 timesteps to 96. Pool sizes were selected by experimentation.

The main processing is done via the Gated recurrent unit (GRU) layer. We chose GRU as related work shows better temporal performance while maintaining equivalent accuracy [85]. Best accuracy was achieved with a unit count of 3 * the number of labels. A Gaussian noise layer was added to prevent overfitting and an RMSprop optimizer was used, with a learning rate of 0.02. Nadam and Adam's optimizers were also tested, but they were worse for prediction accuracy. Learning rates up to 0.03 can be used to improve the speed at the cost of less stable changes between epochs. A learning rate schedule was used to decrease the learning rate over time.

Categorical cross-entropy was used as the loss function and Softmax was used as the activation for the output layer. The model is trained over 50 epochs with a batch size of 32. The batch size can be increased to increase speed, but it will reduce

[1] TensorFlow: https://www.tensorflow.org/

the maximum accuracy that could be achieved and the model will converge slower, requiring more epochs. The epoch count can also be decreased, at the cost of stability. An average accuracy of 93% was recorded, with the validation and training producing similar accuracy. The trained gesture is then loaded as a frozen graph.

For live gesture detection, the capacitive data is filtered and then sent to the model by using TensorFlowSharp [2]. TensorFlowSharp is a runtime that allows for TensorFlow models to run from C# and therefore in Unity. Two 100-sample rolling windows were used for continuous detection, one with 100 latest samples and the other with 100 samples from the previous window. Once a gesture is detected, a Unity event is fired with the corresponding gesture name. Finally, the toolkit source code and 3D models of the hardware components are entirely open source [3].

4.3.3.3 Gesture Detection Accuracy

To test gesture detection accuracy, we conducted a small preliminary study involving 4 participants. Participants would test 3 pre-trained gestures, double-tap finger, double-tap hand, and finger swipe, and a custom gesture created and trained themselves (for a total of 40 recordings). The entire study was done using the 2-sensor radial face configuration. Each of the 4 gestures was tested 20 times by each participant. Detection accuracy for the double-tap finger was 93%, the double-tap hand was 100%, the finger swipe was 85%, and the custom gesture was 85%.

4.4 Applications of Physically-Modifiable Tools

Reflecting on **RQ2.2**, the proposed TangibleTouch toolkit provides a modular and scalable fabrication method for producing physically-modifiable cubic tools, with interchangeable capacitive faces that support a range of surface-based gestures. In this section we showcase a number of interactions using a cubic tool designed with the toolkit, ultimately showing how the interchangeable faces can be used to physically modify the same form factor to offer new functionalities across different applications. We employ a Type 1 evaluation strategy [226] to demonstrate the feasibility of the toolkit and its ability to rapidly prototype a cubic tool with surface gesture interactions. We created three exemplar applications developed in Unity and deployed them across 3 different output spaces: Model Inspector (Augmented Reality), 2D Platformer (Desktop), and a Media Player (Public display). Figure 4.10 shows the generative breadth of the TangibleTouch toolkit, illustrating different touch gestures used in each application.

[2]TensorFlowSharp: https://github.com/migueldeicaza/TensorFlowSharp
[3]TangibleTouch Toolkit: https://github.com/TangibleTouch/Toolkit

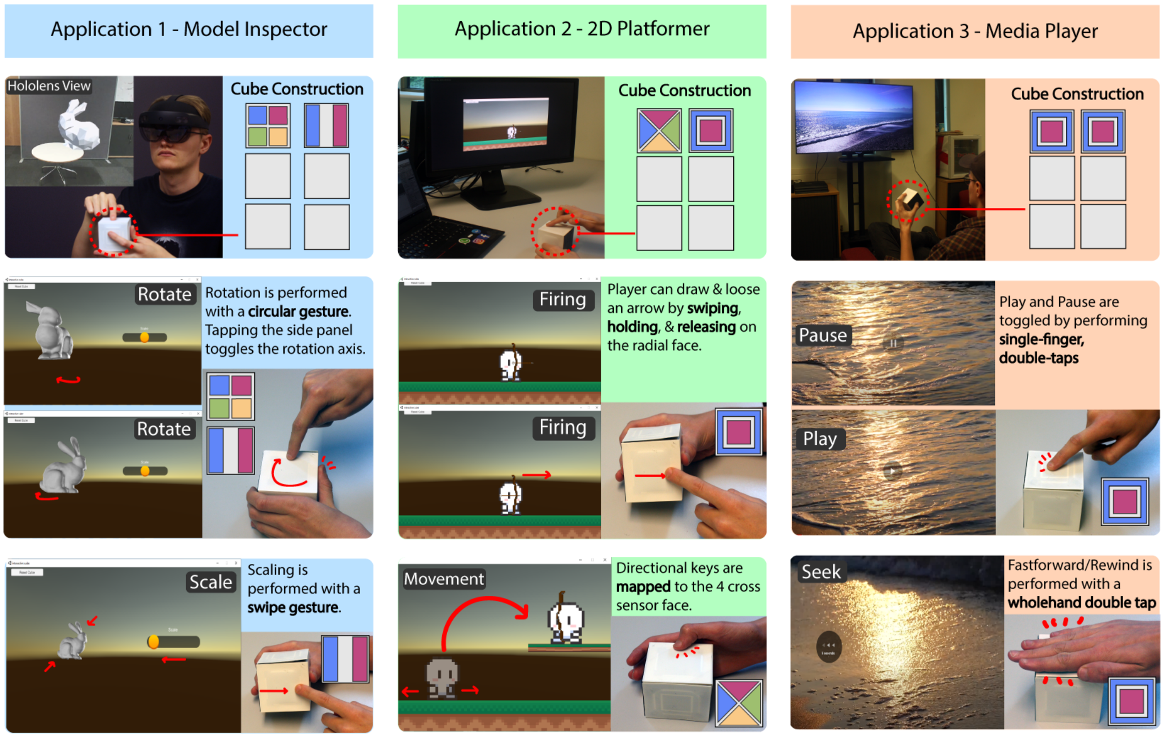

Figure 4.10: Left: Model Inspector application in AR, Middle: 2D Platformer application on a desktop PC, Right: Media Player application on a public display.

4.4.1 Application 1: Model Inspector

This application allows the cube to manipulate a 3D model, loaded into Unity, via rotation and scaling. The cube designed for this application uses 4 blank faces and 2 interactive faces: a 4-sensor square, and a 2-sensor rectangle. The model inspector was deployed in AR, using a Microsoft Hololens2 as shown in Figure 4.10(Left). A 3D model is rotated by performing a circular path gesture on the cube's 4-sensor square face, with the direction of rotation mapped to the direction of the gesture performed. The user can cycle through the 3 different axes for rotation, roll, pitch, and yaw, by tapping either of the two sensors on the rectangle face. Object scaling is performed by swiping from one sensor on the rectangle face to the other, and the direction of the swipe determines whether the object grows or shrinks. The Model Inspector allows users to manipulate and separate the rotational degrees of freedom of a virtual model over distance.

4.4.2 Application 2: 2D Platformer

This application demonstrates a simple platformer game controlling a 2D character to jump, move, draw a bow, and release an arrow. The cube uses 2 interactive faces and 4 blank faces: a 4-sensor cross, and a 2-sensor radial. The application was deployed to a desktop PC shown in Figure 4.10(Middle). In this application, we make use of the key mapping function in the toolkit software to map touch input on the cube's 4-sensor square face to key bindings in Unity, W, A, S, and D, for character movement. The user can perform an additional action of drawing the bow by swiping in any direction on the 2-sensor radial face and then firing the bow by releasing the finger from the sensor.

4.4.3 Application 3: Media Player

The final application demonstrates a simple media controller for playing, pausing, forwarding, and rewinding a video deployed on a public display. The cube designed for this application makes use of 4 blank faces and 2 interactive faces: two 2-sensor radials. This application also makes use of the key mapping function to work with web-based video players. As shown in Figure 4.10(Right), a user can double-tap the centre of one radial face with a single finger to toggle pause and play and perform a whole-hand double tap to fast forward. The same whole-hand gesture is performed on the other radial face to rewind a video.

4.5 Expanding the Toolkit for 3D Tracking

Revisiting the research questions for this chapter, we have addressed the design challenges and opportunities of modifiable tools in AR (**RQ2.1**) and have produced a toolkit that partially addresses the technical challenges (**RQ2.2**). To fully realise the design of cubic tools in AR, there are a series of technical challenges regarding 3D spatial object tracking and user touch input for which we need to expand the *TangibleTouch* toolkit.

One of the most common methods of tracking objects and detecting user interaction is computer vision, for which both marker and marker-less approaches are used in AR [386, 301]. However both approaches have issues with occlusion, especially in handling objects. Other work has tracked objects via sensor instrumentation [183, 192], but this can be problematic for spatial tracking as IMUs are susceptible to drift. Alternatively, object tracking has been achieved using external tracking setups such as lighthouse tracking [277] and infrared trackers which are commercially available but are restrictive in terms of their integration with other form factors, such as cubes, and are also subject to occlusion. More recent work has combined multiple tracking

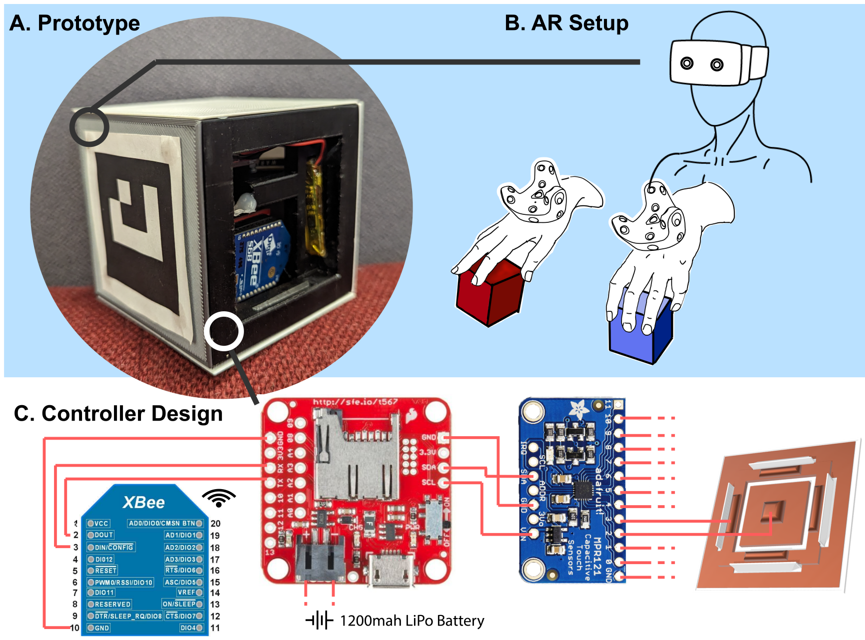

Figure 4.11: The expanded *TangibleTouch* toolkit to incorporate 3D spatial tracking using computer vision for calibration, AR hand-tracking, and onboard IMU in the cubic tool. **A)** Shows the new Cubic tool prototype with a side panel removed and a fiducial marker for calibration. **B)** The AR setup, in this case, we use a video see-through AR display and lighthouse-based hand trackers. **C)** The onboard micro-controller design includes a Razor 9DoF IMU (SAMD21 microprocessor), an XBEE 56B WiFi module, Adafruit MPR121 capacitive sensor, and a 1200mah LiPo battery.

modalities to overcome some of the disadvantages of individual approaches [183, 396]. We incorporate a fusion of these tracking approaches in order to appropriately track interactive cubes in 3D and explore them for interaction. Our tracking system consists of three parts: 1) An AR HMD with computer vision capabilities, 2) a set of bespoke instrumented cubes, and 3) user hand tracking.

For the AR HMD, we used an HTC Vive Pro Eye VR headset combined with a Zed Mini camera mounted to the front of the headset to provide a video see-through display and computer vision capabilities. The camera resolution was 1920x1080 with a vertical FOV of 54°and a horizontal FOV of 85°sampled at 60hz. For the cubes, we adapted the *TangibleTouch* toolkit [310] to incorporate additional sensors to enable full 3D tracking of cubes without relying on computer vision during interactions. This involved replacing the microcontroller with a Razor 9DoF IMU (SAMD21 microprocessor), adding an XBEE 56B WiFi module, and increasing battery capacity to 1200mah LiPo. We used the same Adafruit MPR121 capacitive sensor from the toolkit and opted for the radial face layouts for each face of the cube, resulting in 2 touch sensors per face. We adjusted the toolkit's cube chassis to house the additional components and the resulting cubes measured 6cm^3. The IMU data was sampled at 40hz and sent over WiFi using UDP on a local network to a UDP server within a Unity application simultaneously running the AR environment. Finally, we utilise the HTC Vive lighthouse tracking for the user hand tracking with wrist-mounted HTC Vive trackers which are sampled at 60hz. An overview of the tracking setup is shown in Figure 4.11. While we use this specific tracking setup, the tracking principles can be applied to other AR HMD setups that have hand tracking and computer vision capabilities of which there are a variety [246, 266, 390].

4.5.1 3D Spatial Tracking

The cubes are calibrated using a computer vision marker (OpenCV [376]) which translates the IMU rotational coordinates into Unity coordinates by calculating the difference between the two quaternions and records the baseline values for the IMU and capacitive sensor data (see Figure 4.12). Now we have an understanding of the cube's initial position in space, we rely on a combination of hand tracking and the cubes' IMU to enable 3D tracking, with the main positional tracking achieved by *attaching* and *detaching* the cube to the user's hand trackers as shown in Figure 4.13.

First, we check the cube's last known position to the hand trackers and once within a distance of 20cm we then check the cube's accelerometer data. First, we take the calibrated baseline IMU values and normalise the raw signals to that baseline value. Second, we take the normalised signal of the accelerometer values and apply a basic median filter with a rolling window size of 3, which was sufficient for removing noise, and then apply a OneEuro Filter with a rolling window size equivalent to the IMU frequency (35-40hz). To account for baseline drift in the accelerometer, we calculate a gravity vector for each axis signal and then calculate the magnitude of acceleration. To determine whether a cube should be *attached* to or *detached* from a hand tracker we look at the variance of the magnitude of acceleration over time using a rolling window size of 10. A cube is *attached* to the closest hand tracker if: it is not currently *attached*, it is within the distance threshold to a tracker, at least one of the cube's capacitive surfaces is registering touch input, and the variance of acceleration magnitude exceeds $0.0003 m/s^2$. Likewise, a cube is *detached* from the tracker if: it is currently *attached* if none of the cube's capacitive surfaces are being touched, and if the variance of acceleration magnitude is below $0.00005 m/s^2$.

While this approach allows for full 3D spatial tracking from an initial calibrated position, there are disadvantages to this tracking approach in terms of positional accuracy. Firstly, there is a degree of latency in detecting interactions due to: sending IMU data from the cubes over a WiFi network, processing and filtering the raw data, and applying thresholds to the variance of acceleration. Through testing, we

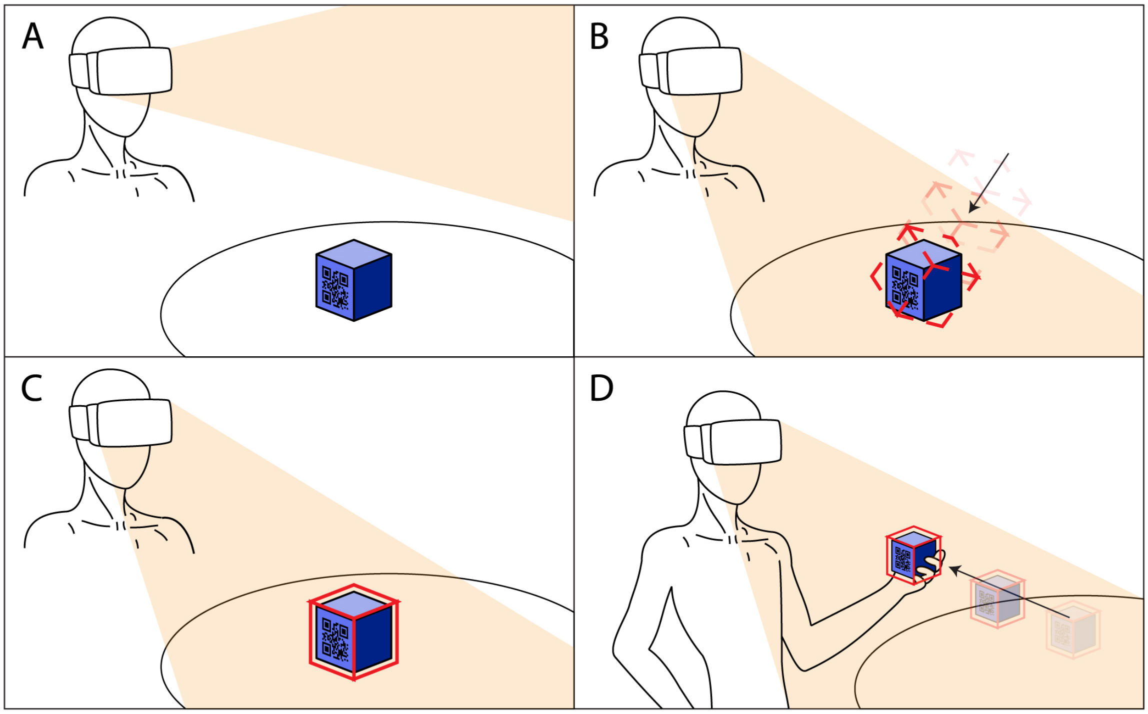

Figure 4.12: The calibration process for enabling 3D tracking of the cubic tools. **A/B)** First, a cube's fiducial marker comes into the user's line of sight and the virtual representation of the cube is snapped to that location. **C)** The virtual cube's rotation is aligned with the physical cube based on the computer vision and IMU data. **D)** The cube is ready to be positionally tracked.

found delays of up to 300(ms) in detecting when a cube should be *attached* and *detached* from a hand tracker, which lead to inaccuracy in positional tracking of up to 10(cm), but this was highly dependent on the movement speed of the user's hand. There is also the question of user grasping and hand pose - which can vary depending on the type of hand tracking employed. Despite the positional inaccuracies we observed, an approximate cube position is always maintained when using the system as positions are recalculated as soon as users interact with the cubes. This is a suitable approach for applications where exact positional tracking is unnecessary but in the cases where improved positional accuracy is required, computer vision can be employed as a corrective measure to re-calibrate cubes 'on-the-fly' as and when their markers are in view of the user. As such, depending on the application different tracking methods can be used more or less without relying entirely on one approach.

Another challenge of this tracking approach is accounting for users moving a cube from one hand to another while it has already been *attached* to a tracker. However,

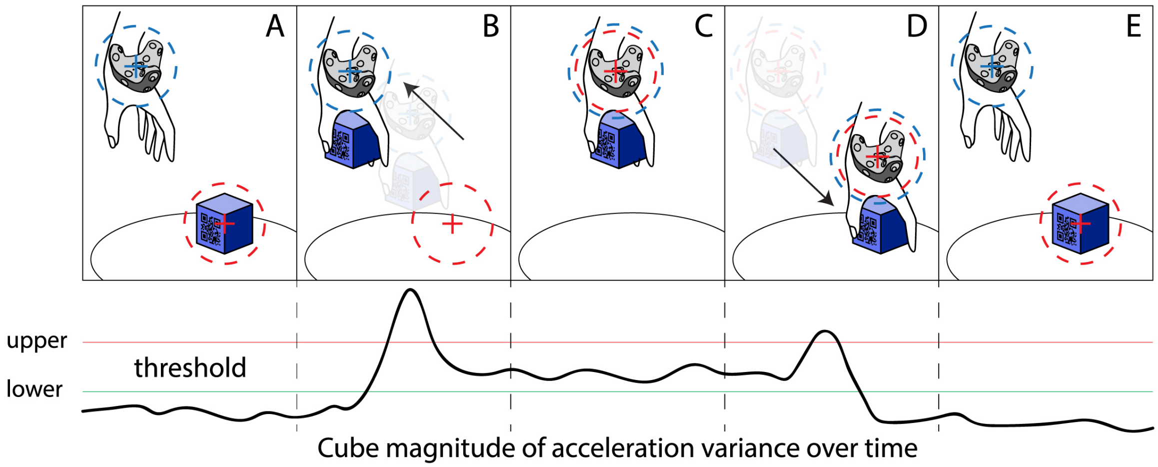

Figure 4.13: The process to attach a cube to a hand tracker. **A)** The hand tracker and cube virtual positions are separate. **B)** The user picks up he cube and the acceleration variance starts to increase. **C)** The acceleration variance of the cube reaches the upper threshold and the cube's position is attached to the hand tracker. **D)** The user places the cube down and begins to move their hand away — the acceleration variance first increases and then decreases as the cube is stationary. **E)** Finally, the acceleration variance drops below the lower threshold, and the cube position is set to the position of a previous frame accounting for the delay in threshold detection.

we use motion correlation between the user's hand acceleration and the acceleration of the cube to disambiguate which hand the user is holding the cube in. For the matching process, we use Pearson's correlation coefficient with the cube and hand tracker acceleration magnitude as input, with a sliding window of 30 frames (0.75 seconds) with the tracker acceleration data offset by 30 frames (0.33 seconds) which gave the best match overall. We then compare the correlation values and *attach* the cube to the hand tracker with the best match. Using motion correlation for *attaching/detaching* the cube was overall less performant than the method described above, so we only employ motion correlation when a cube is already in hand.

4.6 Applications of Virtually-Modifiable Tools

After mapping out the possibilities of cubic tools in AR in Section 4.1, let us operationalise these concepts into concrete examples of AR applications. We developed two demonstrative applications in AR using the cubic tools: *AR Workspace* and *AR Maps*. Each application employs the different interaction metaphors applied to cubic tools (proxies, containers, and handles) described previously.

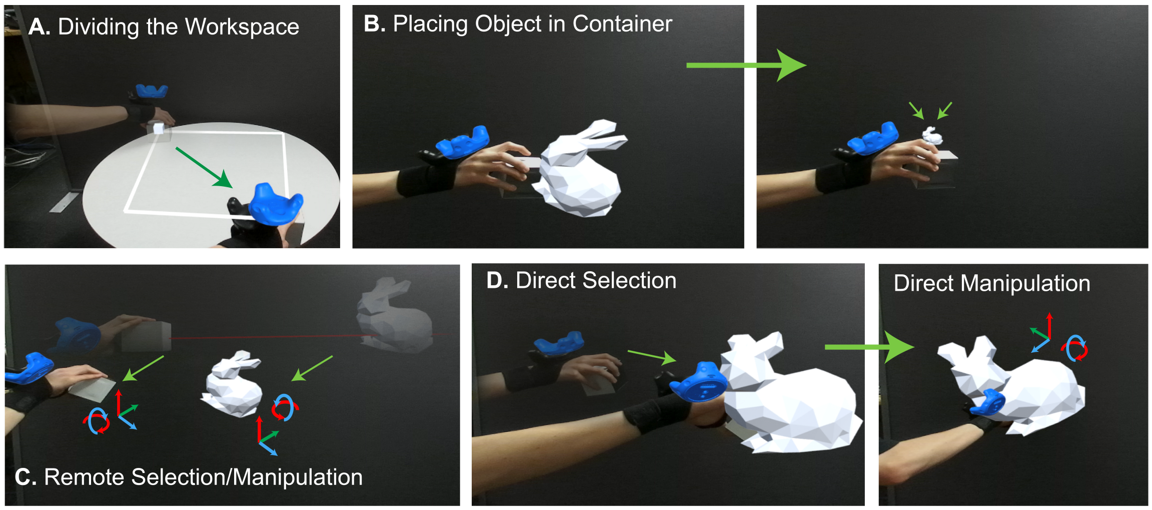

Figure 4.14: The *AR Workspace Application*: **A)** The user divides the workspace by placing an 'area definer' cube, tapping a face, and dragging. **B)** A virtual object is placed in a container by physically colliding a container cube with an object. **C)** A virtual object is 'remotely' selected by a ray from a handle cube intersecting with an object and the user tapping a designated face. The user can then manipulate the object in 6DoF. **D)** 'Direct' selection of a virtual object by physically colliding a handle cube with the object, then providing 6DoF manipulation.

4.6.1 Application 1: AR Workspace

This application integrates proxy, handle, and container cubic tools to enable the seamless blend of a physical and digital workspace. Using the cubic tools the virtual workspace can be organised and divided in a similar capacity to how we organise our physical workspaces. The cubic tools can also be used for file handling of virtual objects and scenes, replicating the desktop-like functions of 'copy', 'move', and 'delete'. Additionally, the cubic tools allow for 3D object manipulation, aligning with the current expectations of AR environments.

4.6.1.1 Dividing the Workspace

Due to a cube's ability to be *spatially arranged* [235], the cubic tools can exploit existing physical surfaces in the AR environment as a scaffold to organise and divide the virtual workspace. A handle cubic tool configured as an 'area definer' can detect the presence of a physical surface when it is placed down through the capacitive sensors on the bottom face. Once placed, a user can tap the top capacitive sensor and start defining the virtual area to the desired size by dragging the cube across the surface and then tapping the top sensor once more. In the example in Figure 4.14, the user is creating a 2D rectangular area, but the area could also be configured to be any shape or dimension by changing parameters in the 'area definer' functionality.

Additionally, in Figure 4.14 the user is creating a virtual object staging area in which a proxy cubic tool in this area that collides with a virtual object will *copy* and *select* the virtual object for 3D manipulation. However, we can imagine that an 'area definer' can be configured to create different types of areas in which the physical cubic tools and virtual objects can behave differently on entering. For example, a modelling area in which a proxy cube is instead used to manipulate the vertices of a virtual object's mesh instead of the entire model, or a display area in which a handle cubic tool placed there can be operated as a slider to rotate a virtual model.

4.6.1.2 Manipulating Virtual Objects

A cube's ability to be *manipulated* [235] is utilised for moving and modelling virtual objects in AR workspace. A cubic tool configured as a proxy in the AR workspace provides basic interactions such as 3D selection and manipulation (shown in Figure4.14), operating similarly to how cube proxies are described in Section 4.1.1. A proxy cubic tool can select a virtual object either directly, by intersecting the cube with the desired object, or remotely, by first toggling a raycast pointer by holding the top face sensor for 2 seconds and then selecting the first virtual object intersecting with the ray by tapping the same top face sensor. Once selected, the proxy cubic tool can then manipulate the virtual object in terms of translation and rotation either in an 'exact' or 'relational' manner depending on the selection. These fundamental functions of selection and manipulation that a proxy cubic tool provides can be expanded to incorporate other functions such as multiple object selection and manipulation, in which aspects of the 'area definer' can be leveraged to select many virtual objects simultaneously.

4.6.1.3 File Handling

A cube's ability to be a *container* and *pedestal for output* [235] is leveraged for basic file handling in AR such as storing and instantiating virtual objects and scenes. Virtual object(s) stored within a cubic tool container is represented as a 3D thumbnail projected above the cube as shown in Figure 4.14. A user can *cycle* through the items stored within a container by swiping from left to right on the cube's top face sensor. An item stored within a container, whether it be a single virtual object or a collection of them, can be *instantiated* into the scene by tapping on a designated side-face sensor on the cube. An item can be instantiated into the scene whilst remaining in the container any number of times. Virtual objects in the AR environment can be *stored* within a container by intersecting the cube with a desired virtual object and tapping the side-face sensor.

Furthermore, multiple virtual objects and even the entire scene, including defined virtual areas, can be stored by placing the container on a physical surface and holding both the top face and designated side face sensors, essentially ‘pulling’ the entire scene within the container. As a first example of container cubic tools, file handling in AR Workspace could be expanded further to incorporate additional functions of containers described in Section 4.1.2. For example, supporting other types of data (documents, data files, audio, video, etc.), deleting items from a container, copying items to other containers, or combining containers together to operate as data pipelines.

4.6.2 Application 2: AR Maps

This application also leverages the interaction metaphors of cubic tools, similar to AR workspace, but with a greater emphasis on surface-supported work. It provides users with a virtual map that can be created on a physical surface and manipulated using cubic tools. The cubic tools also enable 3D street views of the 2D map, marker placement, route sketching, distance measuring, and map file handling. With the ability to save maps into container cubes and load them onto a physical surface, the AR maps application offers an intuitive way to manage maps and explore street views in a 3D environment entirely facilitated by cubic tools (shown in Figure 4.15).

4.6.2.1 Map Controls

We expand the concept of using physical surfaces to scaffold virtual workspaces introduced in AR Workspace further by primarily utilizing a physical surface as the projection space for an AR map. The cube’s ability to be *spatially arranged* [235] serves as the primary means of controlling the map. To create a map, a container cube with a stored map object can be placed on any physical surface. The user taps the top face sensor and adjusts the map size by dragging the cube across the surface, similar to the ‘area definer’ handle in AR Workspace. Panning the map is done by using a handle cubic tool assigned as a ‘map controller’ - the cube is placed on the map directly, which is detected by the bottom face sensor, and dragged across the map to navigate to a desired location. Clutching the map control can be achieved by simply lifting the cube off the physical surface.

To zoom in and out of the map, a second ‘map controller’ cube is introduced on the physical surface. A common input mapping for scaling and zooming is varying the distance between input points [264], using the metaphor of ‘stretching a piece of rubber’ [428]. Therefore, to zoom in, the cubes are simultaneously moved away from each other, and to zoom out, the cubes are simultaneously moved toward each other. The zooming function can be disabled by simply lifting the second cube out of the map space. The map can also be explored in a street-view mode by reassigning the ‘map controller’ cube as a ‘street viewer‘ by swiping a designated side face sensor,

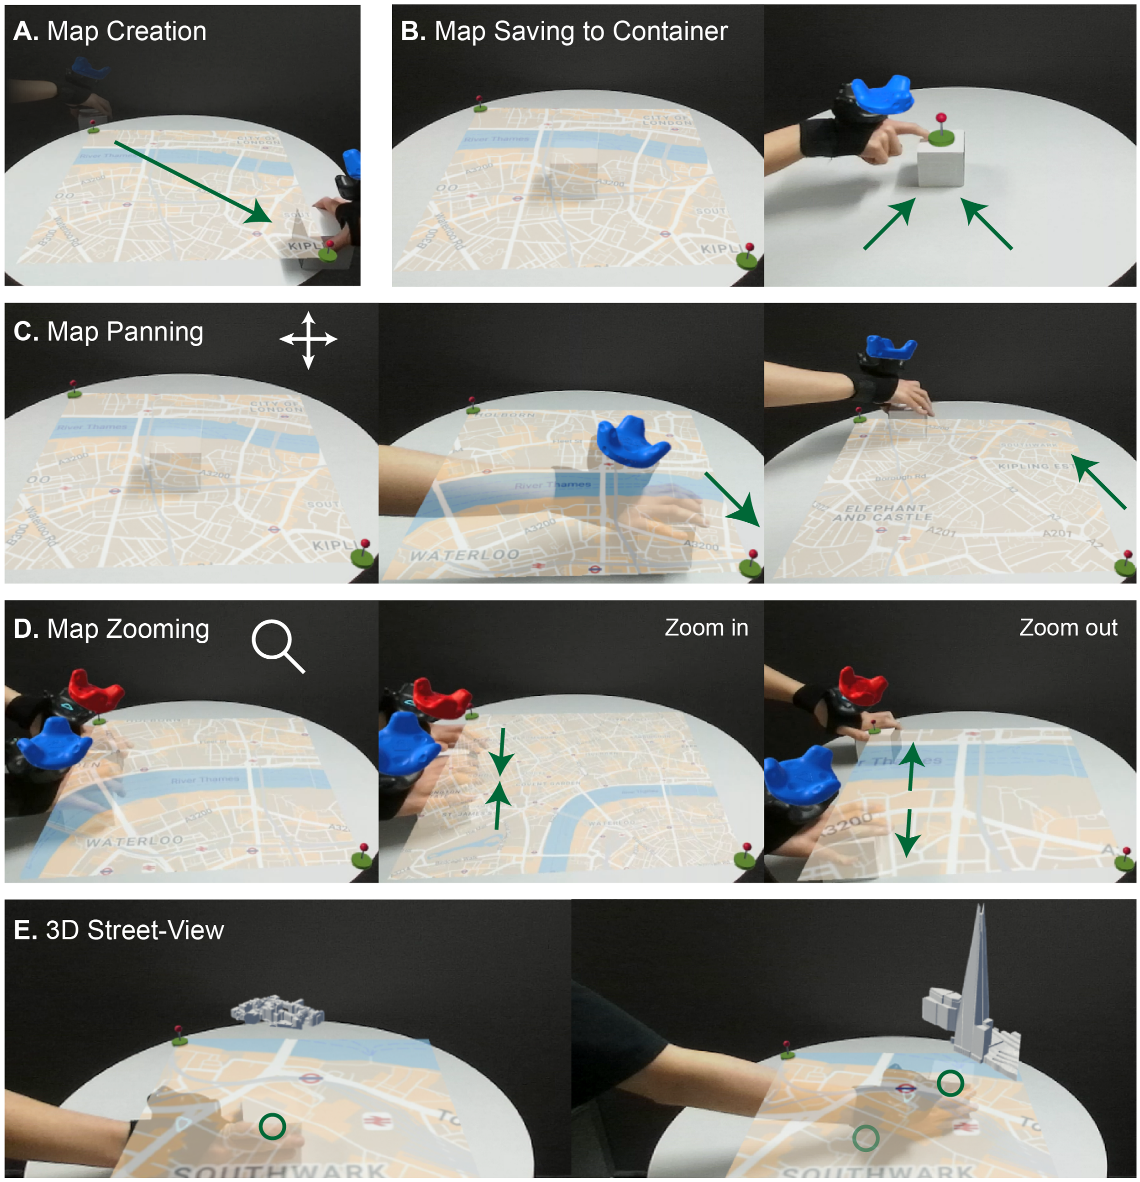

Figure 4.15: The *AR Maps Application*: **A)** A user creates a map from a cube container by placing it on a physical surface, tapping the top face sensor, and dragging it to the desired size. **B)** The user saves a map to a cube container by first placing it on the map, and then tapping the top face sensor which saves the map, represented as a 3D thumbnail. **C)** Map panning using a 'map controller' handle. Placing the cube down and dragging pans the map from that point. **D)** The user introduces a second 'map controller' on the physical surface to perform zooming, moving the cubes towards zooms out the map, and away zooms in. **E)** The user uses a 'street-viewer' handle cube which projects 3D models of the 2D map that the cube is currently on.

Figure 4.16: The *AR Maps Application* (continued): **F)** The user uses a 'marker placer' cube to place down markers on points of interest by tapping a designated cube face. **G)** The user uses a 'route planner' cube in a similar manner to the 'marker placer' except the markers are connected by a route. **H)** The user uses a 'distance measure' cube to first place down an anchor point on the map, by tapping a designated face sensor, and then drags the cube to a desired location. The distance measure is shown above the anchor point.

which cycles the different map functions. The 'street viewer' cube can then be placed directly on the physical surface, and a 3D model of the corresponding portion of the map will be projected above it, leveraging the cube's affordance as a *pedestal* [235]. The street view can be explored by dragging the cube around the 2D map.

In addition to the map controls shown in Figure 4.15, there are various other functions that can be added. For example, map rotation could be achieved by using a cube as a 'dial' on the physical surface. Similar to the AR Workspace 'area definer', different map shapes, sizes, and types can also be explored, such as topographical maps. The projected WIM model using the 'street viewer' cubic tool can also be treated as a manipulable virtual object. For instance, a proxy cube can be used to select the WIM model projected above the 'street viewer' cube and then lifted out of the map space, allowing the user to scale it and inspect it more closely.

4.6.2.2 Markers, Routes, & Distance

We further leverage a cube's affordance to be *spatially arranged* [235] by allowing users to annotate an AR map with markers, routes, and distances — shown in Figure 4.16. To do so, a handle cubic tool can be assigned as a 'marker placer' by cycling through the different map functions via a swipe on a designated side face sensor. Once assigned, the user can place the cube on a physical surface and tap the top face sensor to place a marker on the cube's current location, with the ability to place multiple markers on a single map. To delete a marker, the user simply taps the top face sensor of the 'marker placer' cube when it intersects with an existing marker. Although this function operates similarly to cubic containers, it is conceptually different as the markers themselves are not the primary objects of interest, but rather annotations on an object of interest. They cannot be directly manipulated, and as such, the 'marker placer' cube is classified as a handle, although this functionality could also be designed from the perspective of a container.

To further enhance the functionality of AR maps, we have a 'route planner' tool that allows users to create custom routes by placing numbered markers on the map. A cubic tool can be operated as a 'route planner' first by cycling to that function as described previously. Once assigned, the 'route planner' cube operates similarly to the 'marker placer' cube, but instead of placing markers, it creates numbered markers that are connected by a line, forming the route. To place a route marker, the user simply places the cube on the physical surface and taps the top face sensor. If the cube intersects with an existing route marker, it will delete it and recalculate the route accordingly. For example, if a route exists between markers 1, 2, and 3, and marker 2 is removed, the route will be reformed to connect markers 1 and 3, and the markers will be renumbered to reflect the new route.

Lastly, a cubic tool can be assigned as a 'distance measure' by cycling to that function as described previously. Once assigned, a user can place the 'distance measure' cube on the AR map on the physical surface and tap the top face sensor to establish an anchor point. The cube can then be dragged around the map surface, and a measurement line will be projected between the anchor and the cube, displaying the calculated distance above the anchor. To make the measurement line permanent, a user taps the top face sensor again. To remove it, a user intersects the cube with the anchor and taps the top face sensor once more.

4.6.2.3 File Handling

Finally, an AR map can be saved and loaded, including any annotation markers or routes, in a similar manner to the containers described previously in the AR workspace application and in Section 4.1.2. To load a saved map, a user places the container onto a physical surface, taps the top face sensor, and drags the cube across the surface to the desired map size. To save a map, a user intersects the container cube with the map and taps the top face sensor. This will create a new item within the container and remove the map from the scene.

4.7 Discussion

Prototyping tangible objects with touch capabilities is complex, and designing for intricate surface-based gestures is non-trivial. Additionally, 3D tracking of physical objects in AR without computer vision is difficult to achieve and requires objects to be instrumented. Our toolkit, TangibleTouch, addresses these challenges by providing an extendable and modular hardware platform, that leverages cube affordance, and a software platform for configuring and designing tangible gesture interfaces across many different output spaces. Designers using the toolkit can: i) design their own bespoke sensor configurations for a cube, ii) create, train, and test in real-time a machine-learning model for a set of surface gestures using the provided interface components, and iii) deploy interactions to a variety of applications built in Unity.

We demonstrated how the toolkit can be expanded to include additional sensors and microcontrollers (such as IMUs and WiFi Modules) to improve sensing capabilities. Additionally, our 3D tracking approach of leveraging a combination of AR HMD hand-tracking, computer vision, and instrumented physical objects can be extended to include other form factors and types of objects. Future work can potentially look to reducing the required space for the sensors so that even everyday objects could be instrumented for tracking to further expand on opportunistic controls [156].

Focussing on the aspects of the toolkit for designing surface gestures, expert designers can build upon the capacitive face design we introduced in the hardware platform, to build entirely new face configurations for different or more complex surface gestures. The *TangibleTouch* interface and cube configuration software can support any number of face designs, and gestures can be trained with any sensor configuration. To ensure accessible fabrication, we designed the toolkit components to be produced entirely using single extrusion 3D printing.

Additionally, we show the toolkit's ability to rapidly prototype using accessible fabrication methods. Designers can multiplex surface gestures in a single artefact by simply using the discrete faces inherent to the cubic form factor. The modular nature of the cube allows for on-the-fly reconfiguration of the interaction device, ideal for exploring a breadth of interactions during the prototyping stages of a tangible interface. Also, the abstract nature of cubes means that any designed gestures can often be transferrable to other, more complex tangible form factors.

Using three demonstrative applications that include a number of different sensor configurations and surface-based gestures, we highlight the generality of interactions supported through physically-modifiable tools. The two AR applications (AR-Workspace and AR-Maps) demonstrate that by using the toolkit, the interaction metaphors of virtually-modifiable cubic tools can be realised in AR. However, tool modification as supported by the toolkit prototypes do not provide mechanisms for physical shape change as part of an interaction, something that we explored in the previous chapter. Instead, the cubic tools we developed are modified between interactions and applications, allowing a cube to be repurposed to detect different surface gestures while maintaining its shape. As a result, the concept of tool-modifying is more complex than first anticipated, and we explore this nuance of physically-modifiable tools further in Chapter 6 (Section 6.2.1).

4.7.1 Building upon TangibleTouch

There is a clear avenue for future work building on the foundational elements of TangibleTouch. Firstly, to develop more face designs in terms of capacitive sensor configurations but also explore face surface texture, form, and colour. By using the same modularity principles, there is potential to not only design and explore further surface gestures, but also explore haptic experiences and output. For example, to differentiate and communicate gesture mappings to faces or better convey desired interactions to novice users through the materiality and affordance of a face. To achieve this, more sophisticated fabrication approaches could be incorporated to expand the toolkit design space further.

In terms of leveraging cube affordance in tangible interaction, our toolkit has scratched the surface. The highly decoupled nature of the toolkit enables surface-based gesture detection to be incorporated with additional sensing approaches or devices such as proximity sensors and visual displays. Furthermore, our initial characterisation of the interaction space briefly touched on the implications of multiple cubes. The interplay between multiple interactive cubes and their utility in collaborative tasks warrants exploration in and of itself. The cube form factor could be condensed, by adjusting the level of instrumentation, opening up a design space for 100s of stackable, miniaturised cubes that can be configured into new and unique geometries. Finally, while we have evaluated TangibleTouch through demonstrative applications [226], future work could employ different evaluation methodologies such as case studies, a usability study, or heuristic evaluation.

4.8 Chapter Conclusion

In this chapter we explored the design of physically and virtually modifiable tools for AR focussing on one specific form factor — **cubes**. We answered **RQ2** and the sub-questions by first mapping out the design space for physically and modifiable cubic tools in AR and devised a set of *interaction metaphors* – proxy, container, and handle. Accompanying this we described the potential for interaction when combining a cubic tool and user touch gestures. Using these design spaces as a road map, we then developed and expanded upon the *TangibleTouch* toolkit, enabling the fabrication of instrumented and modifiable cubic tools. Prototypes developed by the toolkit are capable of detecting a wide array of surface gestures. We also presented a novel 3D tracking approach using instrumented cubes, AR HMD hand tracking, and computer vision, something which can be applied to any instrumented object in the future. From this chapter, there are a number of conclusions that can be drawn and opportunities for further exploration:

Firstly, by using our *TangibleTouch* toolkit designers can prototype and develop bespoke surface-based gestures for tangible interfaces using a modular and easily fabricated hardware platform. The provided software framework also abstracts away complex data processing and machine learning and reduces testing complexity by using a run-time environment to detect designed gestures in real time. Both the hardware and software platforms are highly decoupled and extendable for expert designers to create other capacitive face designs, incorporate additional sensors, and implement them into any Unity-based applications — something which we demonstrated by incorporating 3D tracking.

Secondly, our design space exploration of tools in AR and the applications presented in this chapter give a broad overview of what interactions are possible with the cube form factor, how a tool can be physically modified for different interactions, how different *interaction metaphors* can be applied to virtually modify a tool, and how the physical environment can be leveraged for support. Despite the examples presented, there are a number of ways specific interactions can be designed, for instance, 3D manipulation using gain and DoF control using the cubes, beyond what we have already explored. Likewise, it is unclear if certain interaction designs are better suited for the cubic tools over other designs.

Thirdly, while cubic tools have a broad range of interactions that cannot be achieved using current state-of-the-art AR tools (hand-controllers), it is unclear if cubic tools can perform better, the same, or worse in certain canonical AR activities such as 3D manipulation. A clear next step is to comparatively evaluate the cubic tools against pre-existing interaction techniques.

Finally, we found through the development of our toolkit that physically-modifiable tools can be designed and interpreted in different ways. For example, in chapter 3 we specifically focussed on physically-modifiable tools that shape-change *during* an interaction, and instead the toolkit produces physically-modifiable tools that are designed to be reconfigured *between* interactions while maintaining the same form factor. This is something that is discussed further in Chapter 6 (Section 6.2.1).

In the next chapter, we use the prototypes we developed in this chapter to probe variations of a specific interaction technique for precise 3D object manipulation that we designed using cubic tools. The technique itself combines the handle and proxy interaction metaphors discussed in this chapter and utilises some of the key benefits of using cubic tools such as leveraging physical surfaces, *regraspability*, and bimanual interaction. We explore variations of this interaction technique through empirical study, analyse interesting behavioural phenomena around user handedness that arise when using the technique, and finally comparatively evaluate the cubic tools against state-of-the-art AR/VR techniques — hand-controllers.

Chapter 5

Tool-Using

In the previous chapter we explored physically and virtually modifiable tools in AR focussing on one form factor — **cubes**. We described different interaction metaphors and surface-based gestures that can be performed with cubic tools in AR and created a fabrication and prototyping toolkit. We then developed a number of applications to demonstrate the validity and generality of the toolkit and showcase the interaction metaphors in AR. After showing the breadth of interactions available for a cubic tool, in this chapter, we explore and evaluate a specific interaction technique for 3D manipulation of virtual objects developed using the toolkit from the previous chapter. The technique itself uses two cubic tools, one operating as a 1-to-1 spatial proxy for a virtual object and the other as a handle operated on a physical surface to control the control-display gain (C/D gain) between the proxy and the virtual object. Theoretically, this allows for more efficient and precise translation, rotation, and scaling of virtual objects in 3D over very small or large distances.

To address the final research question around *Tool-Use* (**RQ3**) *How do newly designed tools compare to existing interaction techniques and methods for Augmented Reality?* we conducted three empirical studies to explore the design and behavioural phenomena surrounding the technique, and the technique's performance to two VR/AR hand-controller techniques. We begin by describing the interaction technique in detail before introducing the first study of the chapter which compares 3 different designs for controlling the C/D gain when using the cube interaction technique. The next study then explores the role and influence of user handedness when operating the interaction technique as, theoretically, a user can control how much work either hand does when performing manipulations by increasing or decreasing the gain factor. The final study then compares the performance of the cube technique for 3D manipulation tasks against two techniques designed for VR/AR hand controllers as a baseline. The tasks themselves are based on virtual object alignment, sometimes referred to as docking tasks [72, 104, 211, 264], and range in manipulation size from very large (meters) to small (centimeters). In the three studies, performance is compared using a combination of task time, accuracy, user workload/effort, and user hand movements. Finally, we end the chapter by reflecting on the study results and discussing the implications of our findings.

5.1 Technique Design and Implementation

To showcase how a simple form factor such as cubes can be used for complex interactions in AR, we developed a bimanual technique using two interactive cubes to precisely manipulate virtual objects in terms of translation, rotation, and scale. Previous work from Buxton and Myers on bimanual manipulation found that users were able to perform tasks in parallel and even split tasks between two hands, demonstrating a significant performance increase over one-handed techniques [63]. Hinckley et al. further expanded on this with labour division across different input modalities controlled by the dominant and non-dominant hands respectively [170]. For precise interactions, providing spatial constraints [155] and altering the control-display (C/D) gain [316] are two popular approaches, with C/D gain showing a decrease in task time and workload while maintaining precision [7, 50, 68, 375]. In particular for virtual environments, some interactions would not even be possible without adjusting the C/D gain [312, 334, 432]. Considering this, we set out to develop a virtual object manipulation technique that could support precise interactions in AR by combining the benefits of bimanual interaction, C/D gain control, and the advantageous affordances of the cubic form factor such as surface stability and one-handed manipulation (*re-graspability*. As such the technique consists of 3 parts:

1. The manipulation cube
2. The configuration cube
3. Control-Display (C/D) gain

The two cubes are designed to be operated simultaneously in separate hands with the *manipulation cube* designed to be used primarily in-hand and the *configuration cube* used in conjunction with a physical surface, such as a desk or table.

5.1.1 Using the Technique

To select a virtual object, the *manipulation cube* emits a ray that can be used to point at the desired object which can then be selected by tapping the top sensor of the *configuration cube*. Once a selection has been made the principle of the technique is that the *manipulation cube* is used to directly manipulate virtual objects, with the

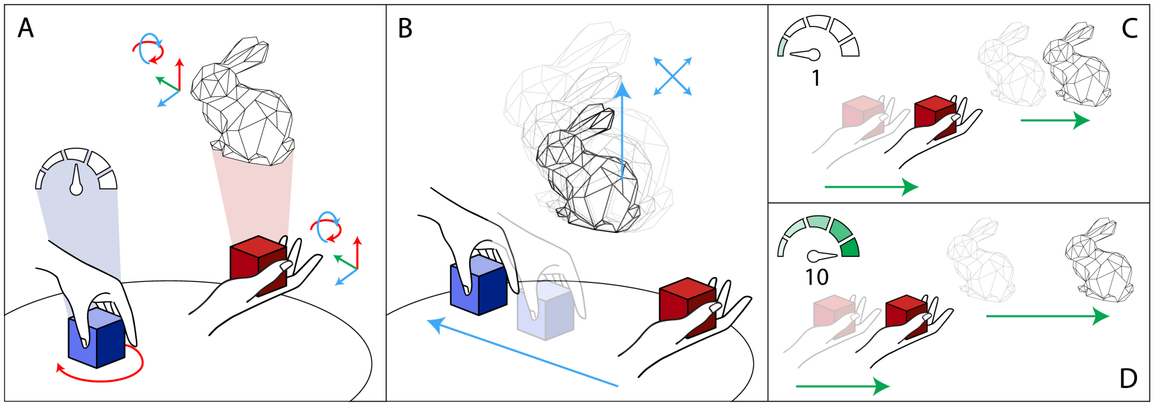

Figure 5.1: The cube interaction technique for precise 3D object manipulation consisting of a *manipulation cube* (MC - red) and a *configuration cube* (CC - blue). **A)** The MC acts as a direct spatial proxy in translation and rotation, while the CC is a rotational dial to increase and decrease the control-display gain between the MC and the virtual object. **B)** The scale of the virtual object is based on the lateral movement of the CC and the C/D gain — moving the CC away from the MC increases the virtual object scale multiplied by the gain, and vice versa moving towards decreases the scale. **C)** Is an example of how the C/D gain affects virtual object translation.

selected object mimicking the translation and rotation of the cube. The virtual object mapping, in this case, is an *exact* mapping, where the movement of the physical cube moves the virtual object a set distance at a constant rate [221]. *Exact* mapping for direct manipulation in VR/AR is a fairly common approach but is limited when moving virtual objects at room or voxel scale. In this case, an algorithm-based mapping can be used to apply a scaling factor between the virtual object and the physical cube to either increase the precision or reach of the user [221]. The Go-Go interaction technique [312] is a classic example of an algorithm-based mapping where the virtual hand of the user moves exponentially faster when the physical hand is extended beyond a predetermined distance threshold.

In the case of our technique, we use the *configuration cube* to explicitly control the gain factor while a user is moving a virtual object with the *manipulation cube* which is shown in Figure 5.1. The gain factor is increased/decreased by rotating the *configuration cube* while it is on a surface either clockwise or anti-clockwise in relation to the user, similar to a dial. A gain value of 1 will result in an exact mapping between the virtual object and *manipulation cube*. Increasing this value will increase the reach of the user with physical movements resulting in larger movements of the virtual object, decreasing this value will increase the precision of the user with physical movements resulting in smaller movements of the virtual object. Users can see the gain value projected on their heads-up display.

The last spatial transformation our technique enables is object scaling, which is achieved by a distance mapping between the *configuration cube* and the *manipulation cube*. Since an exact mapping is not physically possible for scaling, the most common mapping resorts to the variation in the distance between two input points [264], employing the metaphor of 'stretching a piece of rubber' [428]. In the case of our technique, the user moves the *configuration cube* along a physical surface away or towards the *manipulation cube* to increase or decrease the virtual object scale. The same gain factor is used for both translating and scaling a virtual object and, similar to translating an object, a large gain value will increase a user's reach with the movement of the *configuration cube* resulting in larger amounts of scaling/shrinking of the virtual object. Likewise, a small gain value will increase the precision of the user with physical movements resulting in smaller amounts of scaling/shrinking of the virtual object.

Notably, the gain factor affects translating and scaling a virtual object but not rotating. In our pilot testing, we found it much simpler to maintain the direct mapping rather than apply the gain factor for rotation as it was less cognitively overwhelming. Instead, we leverage the cube's ability to be re-grasped in one hand to enable full 360-degree rotation. While using the technique, if the user needs to reposition the *configuration cube* on the surface, to a more comfortable position or to do consecutive scaling, the user can 'clutch' by lifting the *configuration cube* off the surface, move the cube in the air, and place it back down — similar to how a desktop mouse can be 'clutched'. Finally, after performing the manipulations, a user can deselect the virtual object by tapping on the *configuration cube's* top sensor.

5.1.2 Technique Implementation

The direct mapping of the *manipulation cube* to the virtual object for translational movements is achieved using the following equation, where G denotes the gain value, VP is the current virtual object positional vector, and MD is the difference between the manipulation cube current and last frame positional vectors.

$$VP = VP - (MD \times G)$$

To determine whether we should scale a virtual object, we check to see whether the *configuration cube* is moving and in what direction - away or towards the *manipulation cube*. Algorithm 1 demonstrates how we determine whether to scale an object. First, we threshold the Acceleration Variance, AV, of the *configuration cube* to see if the cube is moving. Then, we sum the Total Displacement, TD, of the *configuration cube* over 4 frames by measuring the current Distance to Manipulation Cube, DMC, and the last Distance value, $LDMC$. Next, we threshold the TD to see if the *configuration cube*

Algorithm 1 Scaling a Virtual Object

```
for 4frames do
    if AV >= AVThreshold then
        TD+ = DMC − LDMC
        if TD <= TowardThreshold then
            VOS+ = SF * G
        else if TD >= AwayThreshold then
            VOS− = SF * G
        end if
    end if
end for
```

is moving away or towards the *manipulation cube* to finally increment or decrement the Virtual Object Scale VOS by a Scale Factor SF multiplied by the gain.

5.1.3 Control-Display Gain Factor Design

While we have discussed the mapping between the *manipulation cube* and the virtual object, we also have to consider the mapping of the *configuration cube's* movement to the control of the gain factor. This can be designed in a number of ways and we chose three different gain control designs based on related work and common interaction metaphors. The different designs are shown in Figure 5.2.

5.1.3.1 Absolute

The first gain mapping (Figure 5.2-A) is calculated based on the angular displacement of the *configuration cube* from an original calibrated directional vector. The angular displacement, AD, is measured between a directional vector at the calibrated rotation and the same directional vector from the current rotation of the cube, with this displacement value mapping directly to a gain value, G. $G.MIN$ and $G.MAX$ represents the smallest and largest value the gain factor can be, with $AD.MIN$ and $AD.MAX$ represents the angular displacement values that correspond to $G.MIN$ and $G.MAX$.

$$G = \frac{G.MIN + (G.MAX - G.MIN)}{(AD.MAX - AD.MIN) \times (AD - AD.MIN)}$$

5.1.3.2 Acceleration

The second gain mapping (Figure 5.2-B) is calculated based on the angular velocity of a directional vector from the *configuration cube*. The directional vector's angular

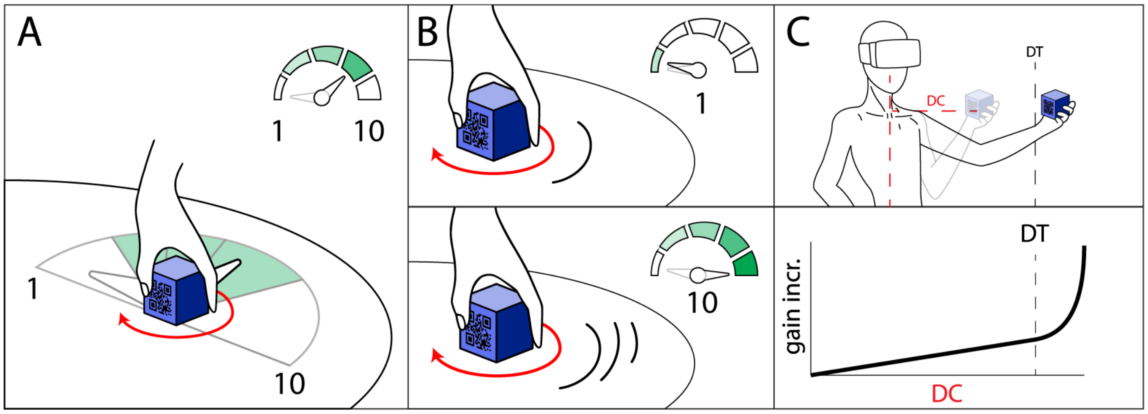

Figure 5.2: The three different gain designs for the cube interaction technique. **A)** *Absolute* gain — the cube's rotation is mapped to a specific gain value. **B** *Acceleration* & *Adaptive* gain — The cube's rotational acceleration is mapped to the gain, the larger the acceleration the bigger the gain increase/decrease. **C** *Adaptive* gain — Inspired by the *Go-Go* interaction technique [312], the distance of the user's hand from their chest multiplies the gain factor. If their hand is beyond the threshold DT the resulting gain value increases exponentially.

velocity, VAV, is measured between the current and previous frame of IMU data, with the angular velocity mapping directly to a gain value, G. Similar to *Absolute* gain, $G.MIN$ and $G.MAX$ represents the smallest and largest value the gain factor can be, with $V.MIN$ and $V.MAX$ represents the vector angular velocity values that correspond to $G.MIN$ and $G.MAX$.

$$G = G + \left(\frac{G.MIN + (G.MAX - G.MIN)}{(V.MAX - V.MIN) \times (VAV - V.MIN)} \right)$$

5.1.3.3 Adaptive

The third gain mapping (Figure 5.2-B & C) is calculated in the exact same manner as the *Acceleration* gain, using the angular velocity of a directional vector VAV from the *configuration cube* to the *Acceleration* gain. The factor, GGF, is calculated using the distance of the *manipulation cube* to the center of the user's chest DC, a distance threshold from the user's chest DT, and a base factor F. The result of this is that the gain increases exponentially as the *manipulation cube* moves further from the distance threshold DT. We designed the adaptive gain this way due to behaviour we observed from users during internal testing and the design of the manipulation technique. To move a virtual object a large distance, a user will first increase the gain rapidly using one hand on the *configuration cube* and then move the *manipulation cube* with the other hand pushing the virtual object as far as they can reach. Then, if they have not reached the target, they will 'clutch' by decreasing the gain down to zero and moving the *manipulation cube* and their arm back to a comfortable position. Finally, they can increase the gain once more to continue moving the virtual object. The key aspect of this gain design is that we assume once a user's arm is stretched out, they will want to freeze the object in place as quickly as possible to more readily return to a comfortable position. As such, once the user's arm is extended over a certain distance threshold, the same rotational movement of the *configuration cube* will increase and decrease the gain more rapidly as opposed to the user's arm being close to their body.

$$G = G + \left(\frac{G.MIN + (G.MAX - G.MIN)}{(V.MAX - V.MIN) \times (VAV - V.MIN)} \right) + GGF$$
$$GGF = DC + (F \times (DC - DT)^2$$

5.2 Study 1: Gain Techniques User Study

To get a better understanding of how accurately and efficiently users can manipulate virtual objects in terms of translation, rotation, and scale we did a preliminary study using the three different gain designs: *absolute*, *acceleration*, and *adaptive*. The study involved users completing a series of virtual object 'docking' tasks [72, 104, 211, 264] where a virtual object is aligned to an identical target virtual object that differs in translation, rotation, and scale. The same virtual model of the Stanford bunny [356] described in Chapter 3 Section 3.3.1 was used for this study to ensure that there were sufficient features on the model for participants to understand the docking/alignment they had to perform by simply looking at their controlled virtual object and the target.

This study was also used as an opportunity to inform our two other study designs (described in Section 5.3 and 5.4), further refining the interaction technique, identifying any interesting participant behaviour using the technique, and elucidating additional AR interface elements to aid task completion. In particular, we want to understand how the 3 gain designs perform compared to one another across a number of *small* and *large* manipulations in terms of the *positional accuracy*, *rotational accuracy*, *scaling accuracy*, and *task time*. Additionally, we quantify the amount of *hand movement* and *hand rotation* for all tasks to understand how different gain designs affect how much users move their hands while using the technique. Finally, we recorded participants' *perceived workload* for each gain design using NASA-TLX.

We recruited 6 participants (4 identified as male and 2 female) with an age of 25-34 (mean: 28.5, SD: 3.15). 4 participants had used VR and AR head-mounted displays occasionally, while 2 participants reported daily use. 5 participants were right-handed and 1 participant was ambidextrous. The study followed a within-subject design with the gain technique, translation size, scale size, and rotation size as factors. For the gain technique there were 3 levels (*absolute, acceleration,* and*adaptive*), 2 levels for target translation size (*small* = 10-99*cm* and *large* = 100-250*cm*), 2 levels for target scale size (*small* = 30-50*cm* and *large* = 70-90*cm*), and 2 levels for target rotation size (*small* = 0-50° and *large* = 70-180°). The chosen values and ranges for translation, scale, and rotation were motivated by a number of factors. Firstly, the *small* and *large* ranges are primarily characterised by being **within** or **beyond** comfortable reaching distance. For example, small and large rotations were configured to be **within** and **beyond** comfortable wrist rotation [73], with small and large translation and scaling configured to be **within** and **beyond** the average human reach [174]. We wanted to first see if the gain function is being leveraged by participants when targets are within and beyond reaching distance.

For rotation specifically, we wanted to observe if participants would re-grasp the cube in hand (only using one hand) while aligning to a target — leveraging one of the key advantages of the device's form factor. Finally, we look at the interaction between the gain technique and the size of manipulation the participants were performing to see if there was an impact on alignment speed or accuracy. The order of conditions was permuted using a counterbalanced Latin square design, with the target values being randomised with their specified ranges. For example, for a **small** scale, rotation, and translation the target could be 50cm distance, 40cm larger, and 45 degrees rotated compared to the control object. Whether a target scale was bigger or smaller than the control object was randomly chosen and the target rotation was altered in all 3 Euler axes. As a result, there were a total of $2 \times 2 \times 2 \times 3 = 24$ permutations with one set of repetitions for a total of 48 trials per participant ($48 \times 6 = 288$ total trials).

5.2.1 Apparatus, Tasks, and Procedure

The study was conducted with the HTC Vive Pro equipped with the Zed Mini video see-through AR camera, 60hz, resolution 2560 x 720, FOV 90° horizontal, 60° vertical, and 100° diagonal max. The study space consisted of a large enclosed tracking space using the HTC Vive Lighthouses and a waist-height table for participants to use the interactive cubes on. Participants remained seated at the table throughout the study and their hands were fitted with two HTC Vive Trackers. Two interactive cubes were used for this portion of the study and were calibrated using fiducial markers to the Vive tracking setup prior to the study, with the capacitive sensor calibrated while the cube is at rest on the table surface. The interactive cubes used in the study are as described in the previous chapter (Section 4.5). Participants were free to decide how to arrange the cubes in their hands before the study, but were asked to maintain their handedness arrangement throughout the study — for example, the *manipulation cube* in the dominant hand. Data from the interactive cube's IMU and capacitive sensors

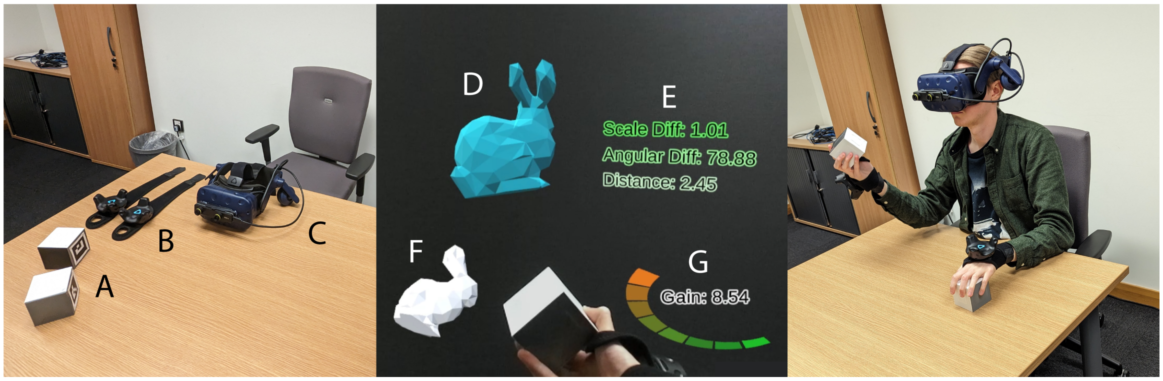

Figure 5.3: The setup used in all three studies. The left image shows the room and equipment setup. **A)** The two cubic tools used to perform the cube 3D manipulation technique. **B)** The Vive hand trackers that the participants wore. **C)** The HTC Vive Pro Eye headset is equipped with the Zed Mini video see-through camera. The middle image shows the participant's perspective during a trial. **D)** The target virtual object the participant hand to align the control object with (**F**). **E)** The visual cue text to aid a participant with the alignment task. **G)** The gain visualiser on the participant's heads-up display. The speed-o-meter visual was used in Study 2 and Study 3, but only the text was shown in Study 1. The right image is a user during a mock trial.

and head and hands position/rotation were recorded at a sample rate of 60hz inside Unity. Before conducting the study with participants, internal pilot testing was first used to decide the parameter values for the different gain techniques. Figure 5.3 shows the study setup along with the participant's perspective whilst wearing the HMD. For all three gain designs the gain minimum value is $G.MIN = 0.01$ and the gain maximum $G.MAX = 18.0$.

Absolute Gain: The angular displacement minimum $AD.MIN = -20.0°$ and the angular displacement maximum $AD.MAX = 270°$.

Acceleration Gain: The velocity of the directional vector minimum $V.MIN = 0.6°/s$ and the maximum $V.MAX = 53°/s$.

Adaptive Gain: The same values for $V.MIN$ and $V.MAX$ are used as in the acceleration gain. For the centre of the chest, DC, we measure from the point 15cm below the HMD position. The value for DT is chosen as $\frac{2}{3}$ of the participant's arm length, typically 30-35cms, and the factor $F = 1$.

To test our three different gain designs, participants performed a series of alignment tasks inspired by similar previous work in AR manipulations [111] and docking tasks [104]. During a task, participants are shown two versions of the same virtual object and instructed to align the grey object (a control object which they manipulate with the cubes) over the blue transparent object (the target). Control objects were always automatically selected for the participant at the start of a trial. We included several cues to aid with alignment and placement, shown in Figure 5.3, that were projected next to the target object displaying the distance, scale difference, and angular distance between the control object and the target. The visual cue text would start to turn green from white once the control object is within a distance of *20cm*, an angular difference of 20°, and a scale difference of *20cm*. For this study, we applied self-defined termination of the task, something which is often applied [38, 76, 111], as we wanted to characterise the ceiling of accuracy. As such participants were instructed to align the control object as close to the target as possible, with a secondary objective to be as quick as possible.

Prior to the study, participants completed a basic demographics form, followed by a 10-minute induction to the study, AR system, and cube interaction techniques. Participants would then perform 3 sessions of trials, 16 trials per session, with 1 session for each gain technique. Before starting a session of trials with a gain technique, participants were given 5-10 minutes of test trials to become accustomed to the interaction and gain control technique. During the test trials, participants were instructed to decide on a particular handedness arrangement — i.e. manipulation cube in their dominant hand and configuration cube in the other — which they had to maintain throughout the actual trials. There were opportunities to take breaks between sessions and between individual trials. During a trial, participants would perform an alignment and instruct the experimenter once they were satisfied with their accuracy to end the trial. Each trial was capped in terms of time to 90 seconds, at which point the trial would automatically end. Upon completion of a session, participants were asked to complete the provided NASA-TLX form and share their thoughts with the experimenter on the gain technique they had just used.

5.2.2 Study 1 Results

The measures for this study were task time, positional accuracy, rotational accuracy, scale accuracy, hand movement and rotation, and perceived workload and usability. The analysis was performed with a four-way repeated measures ANOVA ($\alpha = 0.05$) with Gain Technique, Translation Size, Rotation Size, and Scale Size as independent variables. QQ plots were used to validate the normality of the data. If the assumption of Sphereicity was violated as shown with Mauchly's test, Greenhouse-Geisser corrected values were used for analysis. In the case of significant results, Bonferroni-corrected posthoc tests were used. Effect sizes are reported as partial eta squared (η_p^2). RAW NASA-TLX scores were analysed using Friedman tests and, in the case of significant results, Bonferroni-corrected Wilcoxon signed-rank tests were used

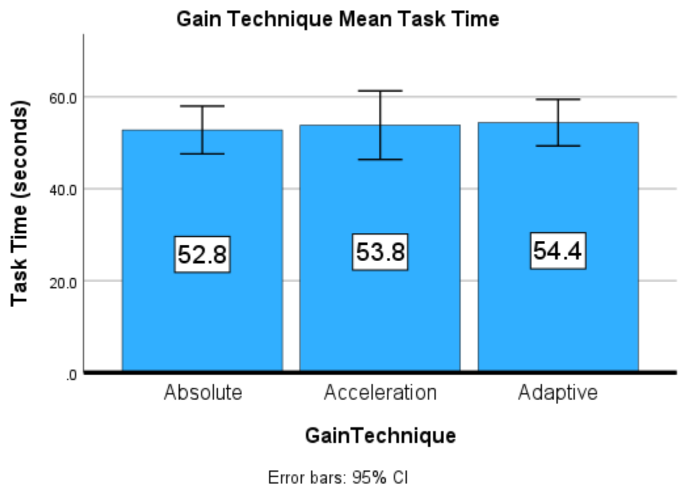

Figure 5.4: The mean task time in seconds of each gain technique. $n = 6$

for posthoc analysis. Temporal filtering was used to remove outlier trials which we define as trials that exceeded the 90-second cap for completion. In total, 40 trials were discarded (13.89%) and after outlier filtering only 3 of the 144 cells had missing values for the four-way repeated measures ANOVA. These missing values were replaced with the maximum value over all participants for each measurement: task time, positional accuracy, rotational accuracy, scale accuracy, hand movement, and hand rotation. Any additional significant outliers identified in SPSS were winsorized.

5.2.2.1 Task Time

To understand how quickly participants could align a control object to a target using the cube interaction technique and to understand if the gain technique influenced the efficiency of alignment, we measured the task time for participants to complete an alignment. Task time is measured from the start of the trial to the moment when participants are satisfied with their alignment and instruct the experimenter to end the trial. The results found no significant 4-way, 3-way, or 2-way interactions between any of the conditions when measuring task time. Likewise, there were no significant differences in task time between the 3 different gain techniques ($F(2, 10) = .120$, $p > .05$, $\eta_p^2 = .024$). **Absolute** gain had the lowest average task time ($M = 52.79$, $SD = 2.02$), followed by **Acceleration** gain ($M = 53.83$, $SD = 2.91$), and **Adaptive** gain ($M = 54.368$, $SD = 1.97$). However, there was a significant main effect of translation size: *small* (10-99cm) and *large* (100-250cm) translation ($F(1, 5) = 7.443$, $p < .05$, $\eta_p^2 = .598$). On further inspection, *small* translations were significantly faster ($M = 49.22$, $SD = 0.63$) than *large* translations ($M = 58.1$, $SD = 2.95$). However, the same effect was not observed for different rotation sizes ($F(1, 5) = 1.365$, $p > .05$, $\eta_p^2 = .214$) or scale sizes ($F(1, 5) = .628$, $p > .05$, $\eta_p^2 = .112$).

5.2.2.2 Accuracy

To understand the alignment accuracy participants could achieve using the cube techniques and to understand if the gain design influenced alignment accuracy, we took 3 different measures: positional, rotational, and scale difference. The positional difference is measured as the distance, in centimeters, between the control object's point of origin to the target's point of origin. The rotational difference is measured as the angular difference, in degrees, between the control object's rotation and the target's rotation. Scale difference is measured as the difference in scale magnitude, in centimeters, between the control object and the target.

The results found no significant 4-way, 3-way, or 2-way interactions between any of the conditions when measuring positional differences. Likewise there were no significant main effects in positional difference between the 3 different gain techniques ($F(2, 10) = .203$, $p > .05$, $\eta_p^2 = .039$). **Acceleration** gain and **Adaptive gain** had the lowest average positional difference (Acceleration, $M = 4.2cm$, $SD = .5$- Adaptive, $M = 4.2cm$, $SD = .6$), followed by **Absolute** gain ($M = 4.4cm$, $SD = .7$). However there was a significant main effect for translation size: *small* (10-99cm) and *large* (100-250cm) translation ($F(1, 5) = 14.843$, $p < .05$, $\eta_p^2 = .864$). Participants were significantly more positionally accurate in *small* translations ($M = 3cm$, $SD = .3$) versus *large* translations ($M = 5.5cm$, $SD = .8$).

For rotational difference, the results found no significant 4-way or 3-way interactions between the conditions, but a significant 2-way interaction was found for $GainTechnique \times RotationSize$ ($F(2, 10) = 4.930$, $p < .05$, $\eta_p^2 = .670$). However pairwise comparison showed no significant simple effects between the different gain techniques or rotation sizes. There were no significant main effects for different gain techniques on rotational difference: **absolute** ($M = 6.67°$, $SD = .899$), **acceleration** ($M = 5.82°$, $SD = .791$), and **adaptive** ($M = 5.89°$, $SD = .511$).

For scale accuracy, the results found a significant 4-way interaction between the conditions ($F(2, 10) = 4.995$, $p < .05$, $\eta_p^2 = .500$). On further inspection, there were significant simple effects for scale accuracy for the **Acceleration** and **Adaptive** gain techniques, specifically when comparing different sizes of manipulations. For the **Acceleration** gain *small translations* differences were significantly more accurate when aligning scale than *large translation* differences, and likewise *small rotations* were significantly more accurate than *large rotations*. However, interestingly the inverse was true for rotation size with the **Adaptive** gain, with *large rotation* differences being significantly more accurate than *small rotation* when aligning scale. This was only prevalent for manipulations that had exactly small scale and small

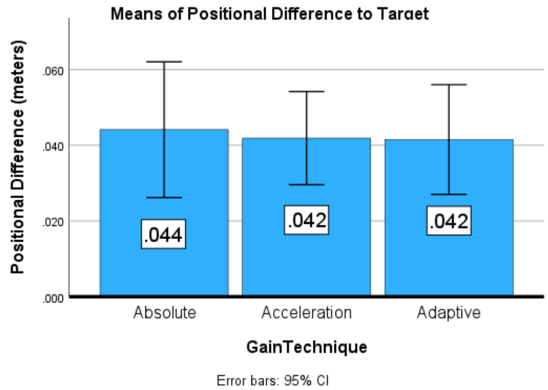

(a) The mean position alignment accuracy of each gain technique in meters.

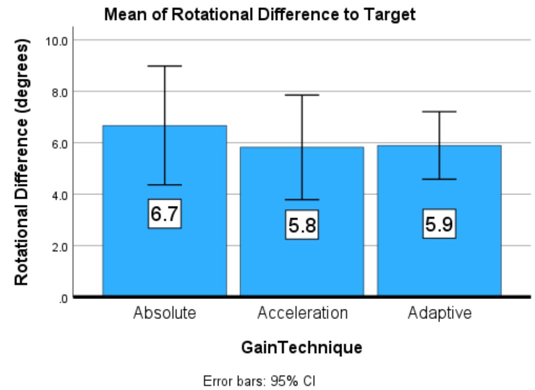

(b) The mean rotation alignment accuracy of each gain technique in degrees.

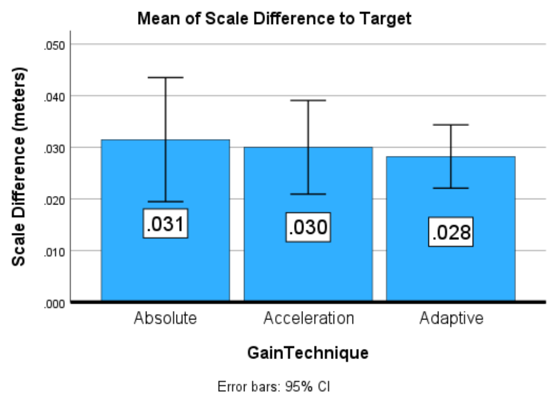

(c) The mean scale alignment accuracy of each gain technique in meters.

Figure 5.5: The mean alignment accuracy of each gain technique for position (Figure 5.5a), rotation (Figure 5.5b), and scale (Figure 5.5c). $n = 6$

translation differences. The results showed no significant main effect in scale accuracy for the 3 different gain techniques ($F(2, 10) = .278$, $p > .05$, $\eta_p^2 = .053$): **Absolute** ($M = 3.1cm$, $SD = .5$), **Acceleration** ($M = 3.0cm$, $SD = .4$), and **Adaptive** ($M = 2.8cm$, $SD = .2$).

5.2.2.3 Movement, Workload, and Usability

To understand how different gain techniques influenced physical movement, user workload, and usability we measured the total physical hand movement and rotation of both hands using the hand-mounted Vive trackers and recorded participant's NASA-TLX scores for each technique. We define total hand movement as the total distance

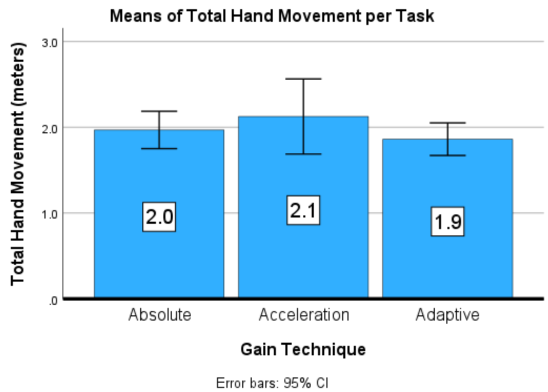

(a) The mean hand movement of each gain technique in meters.

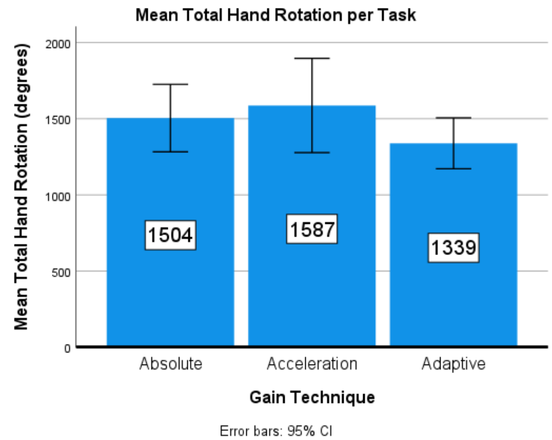

(b) The mean hand rotation of each gain technique in meters.

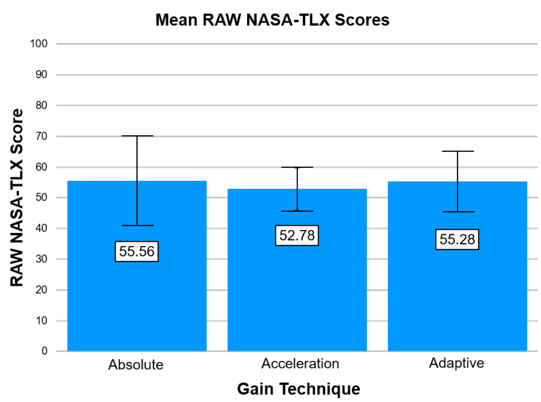

(c) The mean RAW NASA-TLX workload value for each gain technique.

Figure 5.6: The mean user movement and workload. $n = 6$

the hands moved during a trial in meters, and total hand rotation as the tallied rotational movement of both hands during a trial in degrees. While these measures give an indication of physical movement, we stress that these are not objective measures of physical fatigue but rather a measure to compare with the self-reported workload in the NASA-TLX.

For total hand movement, the results found no significant 4-way, 3-way, or 2-way interactions between any of the conditions involving the gain techniques and, likewise, there were no significant main effects between the 3 different techniques ($F(2, 10) = 1.378$, $p > .05$, $\eta_p^2 = .216$). **Acceleration** gain had the highest average hand movement ($M = 2.125m$, $SD = .170$), followed by **Absolute** gain ($M =$

$1.968m$, $SD = .085$), and then **Adaptive** gain ($M = 1.861m$, $SD = .74$). However, there was a significant main effect for different translation sizes ($F(1,5) = 23.999$, $p < .05$, $\eta_p^2 = .828$), with *large* translations involving significantly more hand movement ($M = 2.235m$, $SD = .107$) than *small* translations ($M = 1.734m$, $SD = .067$).

For total hand rotation, we found no significant 4-way, 3-way, or 2-way interactions between any of the conditions involving the gain techniques and, likewise, there were no significant main effects between the 3 different techniques ($F(2,10) = 2.035$, $p > .05$, $\eta_p^2 = .289$). However, there was a significant main effect for different translation sizes ($F(1,5) = 10.686$, $p < .05$, $\eta_p^2 = .681$) and rotation sizes ($F(1,5) = 10.706$, $p < .05$, $\eta_p^2 = .682$). *Large translations* involved significantly more hand rotation ($M = 1641.155°$, $SD = 101.205$) than *small translations* ($M = 1311.728°$, $SD = 42.272$), with the same being observed for *large rotations* ($M = 1546.244°$, $SD = 67.458$) versus *small rotations* ($M = 1406.639°$, $SD = 57.544$).

Friedman tests using the RAW NASA-TLX workload metric revealed no significant differences between the 3 different gain techniques for the overall workload ($\chi^2(2, N = 6) = .333$, $p > .05$). Additionally, there were no significant differences for Mental Demand ($\chi^2(2, N = 6) = .5.778$, $p > .05$), Physical Demand ($\chi^2(2, N = 6) = 2$, $p > .05$), Temporal Demand ($\chi^2(2, N = 6) = .818$, $p > .05$), Performance ($\chi^2(2, N = 6) = 2.7$, $p > .05$), Effort ($\chi^2(2, N = 6) = 1.238$, $p > .05$), or Frustration ($\chi^2(2, N = 6) = .381$, $p > .05$).

5.2.2.4 Trial Observations and Participant Strategy

Based on observations in the trial, we further investigated individual trials to elucidate participants' strategies to the manipulation tasks when using the cube technique as well as differences in operating the different gain techniques. Participants were observed adopting a general approach when using the cube technique to move virtual objects to a target which consisted of planning, moving, refining, reassessing, and correcting. An example of this strategy is shown in a trial by P3 in Figure 5.7a. The four different phases of the strategy are defined as:

- *Planning Phase*: Participants take time to compare the target and the control object and decide on initial movements for position, rotation, and scale.
- *Movement Phase*: Participants enact on their plan to align the two objects. This is characterised by large decreases in the distance, angular difference, and scale difference between the control object and the target. The order and manner in which these movements occur is dependent on the target and the participant's approach. Generally, there are two approaches to the movement — *synchronous* and *asynchronous*, described below. In cases where the control object is already close to the target for a given type of manipulation, the movement phase is skipped and the participant goes straight to refinement.

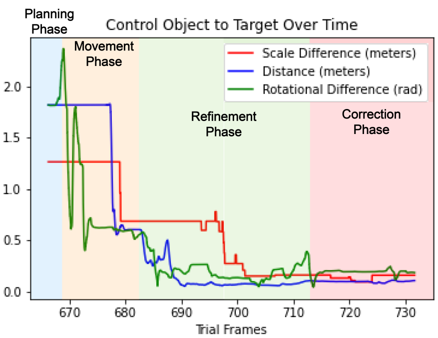

(a) A typical participant strategy (P3) to the alignment tasks. Object manipulation consists of a planning phase, a synchronous movement phase for all 3 manipulations, a segmented refinement phase, and finally a correction phase to complete the alignment.

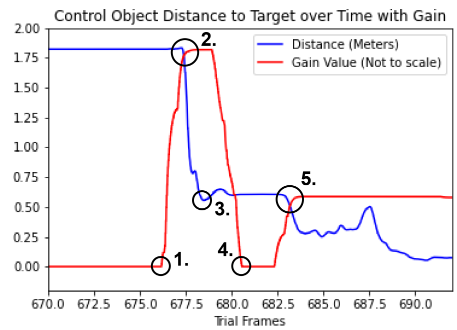

(b) An example trial from P3 of using the gain to clutch the control object while moving. **1)** The user increases the gain rapidly. **2)** The user begins a large translation to the target. **3)** The control object reaches the target. **4)** The user reduces the gain and adjusts their posture for comfort. **5)** The gain is increased once more for refinement.

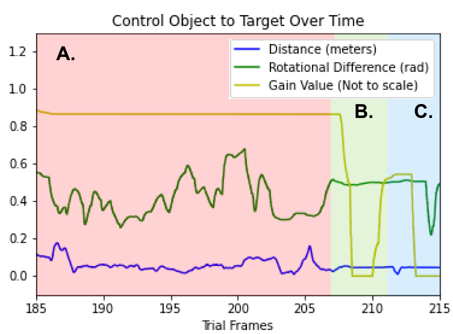

(c) An example of how different manipulations can affect each other. **A)** P1 is attempting to align the rotation of the control object but is affecting the position. **B)** P1 realises this and compensates by reducing the gain to not cause a further change in position. **C)** P1 continues aligning the rotation.

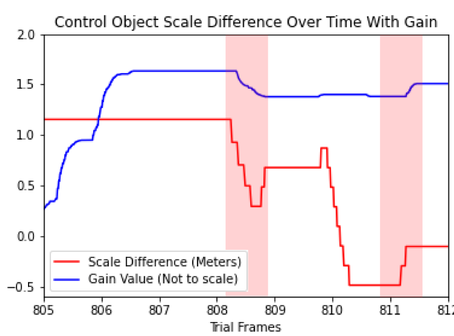

(d) An example from a P5's trial using the *adaptive* gain technique. Changing the scale of the virtual object unintentionally decreases and increases the gain factor.

Figure 5.7: Trial data from different participants. a) shows a typical strategy for the alignment task. b) shows a participant moving and clutching the control object. c) shows one manipulation affecting the accuracy of another. d) shows an example of scaling the control object affecting the gain factor.

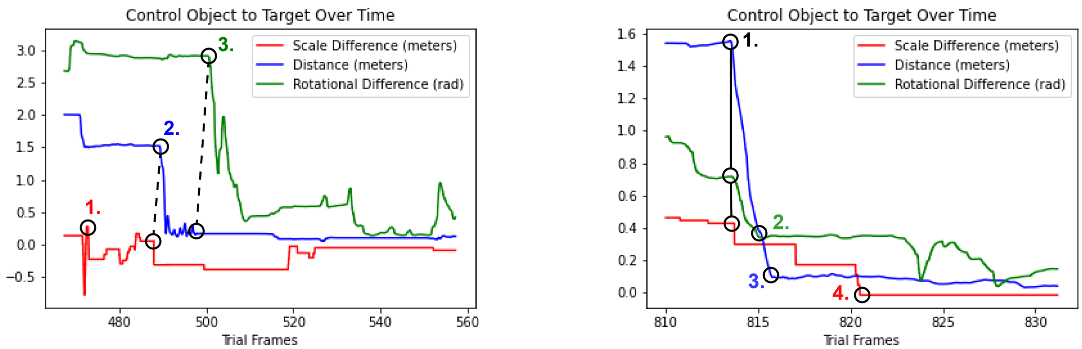

(a) Asynchronous movement of aligning the control object to the target.**1.** The user adjusts the scale to the target size. **2.** The control object is moved toward the target. **3.** The rotations are aligned. This is followed by segmented refinement for first scale and then rotation.

(b) Synchronous movement of the object in translation, rotation, and scale. **1.** Manipulations are started and aligned over time simultaneously, **2.** the rotation is aligned first, **3.** followed by the distance, **4.** and then scale. This is followed by a period of segmented refinement — mostly rotation.

Figure 5.8: Two example trials, P2 (left) and P5 (right), showing Asynchronous and Synchronous movement of a control object using the cube technique.

- *Refinement Phase*: This phase typically succeeds the movement phase after a large set of changes have been made to the control object's position, rotation, and/or scale. Refinement is characterised by small manipulations of the control object over a longer period of time. This phase is usually heavily segmented with participants focussing on one type of manipulation at a time.

- *Correction Phase*: While reassessment occurs throughout the alignment of the control object, the correction phase requires reassessment to recover from errors induced by the refinement phase. Due to the refinement phase being heavily segmented, participants can sometimes induce errors in one type of manipulation while performing another. For example, if trying to refine the object position, a participant may unintentionally rotate the object.

While these were the general behaviours we observed during the study, not every strategy matched with Figure 5.7a. Some Participants would perform *movements* for a manipulation, e.g. translation, then *refine* this before moving on to the next large movement. However, this would sometimes be costly in terms of efficiency as the subsequent large movements would induce errors in previously refined manipulations. Figure 5.7c shows an example of how adjusting rotation impacted translation during a movement if the gain factor is too high.

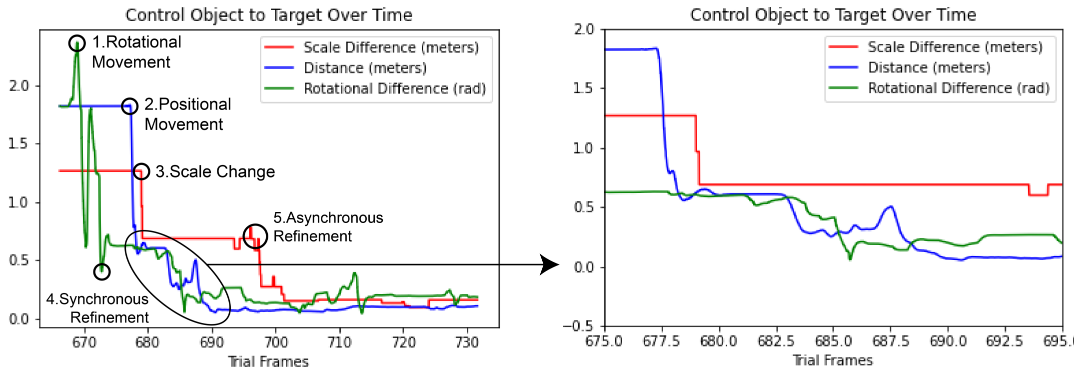

Figure 5.9: Left shows asynchronous movement - (**1.**) rotation, (**2.**) position, and (**3.**) scale. This is then followed by synchronous refinement (**4.**) of position and rotation which is shown in more detail on the right. The user then refines the scale (**5.**)

Participants using the 'clutching' behaviour (shown in Figure 5.7b) were universal. This is a behaviour that came up in testing and something that we designed the *adaptive* gain to streamline. This clutching is characterised by an initial large increase of the gain, followed by a large manipulation of the object in either translation or scale to better match the target. Once the desired position or scale is achieved the gain is then rapidly decreased to 'freeze' the object in place. The user can then readopt a more comfortable position and plan for the refinement step. The gain is then steadily increased slightly while a refinement manipulation is performed. As mentioned, there were different variations of the movement phase observed shown in Figure 5.8:

- *Asynchronous Movement*: Each type of manipulation movement (translation, rotation, and scale) are started and completed before moving on to the next manipulation. Figure 5.8a shows how a scale movement is completed, followed by a positional movement, followed by a rotational movement.

- *Synchronous Movement*: Each manipulation movement is started simultaneously and may/may not be completed simultaneously. Figure 5.8b shows positional, rotational, and scale movement started simultaneously but ended at different times depending on the alignment to the target. The refinement phase is still segmented.

These different movement approaches were not always mutually exclusive. Figure 5.9 shows a participant using a combination of synchronous and asynchronous movement at different times during an alignment. One observation of an unintended interaction prevalent in the *adaptive* and *acceleration* gain techniques is that scaling an object would occasionally alter the gain and vice versa (see Figure 5.7d).

5.2.3 Study 1 Discussion

The key results to inform the subsequent studies are that the expected task time would be between **52-55 seconds**, the expected positional accuracy (approximately) would be **4cm**, the expected rotational accuracy would be **6°**, and the expected scale accuracy would be **3cm**. The study results demonstrate that the different gain technique designs did not have a significant impact on the overall efficiency and accuracy of the cube technique when performing virtual object manipulations.

Reflecting on the significant results, translation size had a significant impact on task time, hand movement, and hand rotation. This implies that while having access to gain control does allow for larger translations more quickly with less effort, it does not equalise the required time and effort of the user between targets that are 'within' and 'outside' of the user's reach. Translation size also impacted positional accuracy, with larger translations being significantly less accurate. This could mean that the visual cues guiding the alignment displayed at the side of the target might need to be improved to equalise positional accuracy when the target is more difficult to see. A counterpoint to this is that translation size did not have a significant effect on scale or rotational accuracy. Instead, this significant difference might indicate that there is decreased positional accuracy as the target distance increases when using the cube technique, which could be accounted for with interface elements such as improved visual cues or object 'snapping'. Interestingly rotation size did not have any significant effect on task time or rotational accuracy which could be a consequence of the cube providing re-graspability — the ability to orientate the cube using one hand. On the other hand, while there was not significantly more hand movement there was significantly more hand rotation for targets with larger rotational differences. This suggests that while re-graspability helps maintain rotational accuracy across different target orientations, it does not equalise the amount of physical hand rotation.

Considering scale accuracy, the *acceleration* and *adaptive* techniques had significantly worse accuracy for certain target conditions, but the same was not true for *absolute* gain. This could be a consequence of the unintended interactions when scaling an object which was observed in some of the trials when using the *acceleration* and *adaptive* gain. Figure 5.7d shows how the gain could change while scaling an object, but also vice versa was observed. In some cases when participants would rapidly increase or decrease the gain, the scale of the virtual object would unintentionally be changed. This could be due to users accidentally moving the cube along the surface while rotating it rapidly, which the system falsely detected as a 'scale' motion. It could be beneficial to decouple the scale and gain functions more for the cube technique, however, this issue was less prevalent in the *absolute* technique. The fact that both *acceleration* and *adaptive* gain rely on the cube's angular acceleration could mean that this unintended interaction is more prevalent using these techniques. As such for *acceleration* and *adaptive* gain, it could be that a more sophisticated implementation is required rather than simply thresholding the cube's angular acceleration.

Another interesting result is that larger target scale sizes did not have a significant impact on task time, accuracy, or physical movement. This implies that the gain control is working as intended for different target scale sizes, equalising the required time and physical movement while maintaining accuracy. However, the range of scale sizes we explored was limited (*30cm* - *90cm*). To better understand the efficacy of the cube technique at scaling virtual objects, more scale sizes could be tested. Furthermore, as per the NASA-TLX results, the gain techniques did not have significantly different perceived workloads. We can infer that the specific design of the gain control does not influence the overall workload of the manipulation tasks. Additionally, the added complexity of the *adaptive* gain technique — utilising the distance of the *manipulation cube* from the user — did not result in higher mental demand but equally did not yield significantly less perceived physical demand. While not significant, it is notable that the *adaptive* gain did have lower mean hand movement than the other two techniques. The 'clutching' behaviour observed during developmental testing was widely observed in the study (Figure 5.7b), which means the *adaptive* design was not unfounded just not significantly beneficial.

Despite the lack of significant results for the three gain techniques, several participants provided verbal feedback on their experience with some of the techniques. P1, P4, and P6 expressed sentiments of the *adaptive* gain being 'too confusing'. P2 mentioned that they "didn't notice the difference" between *acceleration* and *adaptive* gain. Finally, P1 commented that the *acceleration* gain felt 'less precise for controlling the gain' than the other two techniques. Additionally, most participants felt that the current textual representation of the gain factor on the HUD did not convey the gain value or the impact it would have on the control object clearly enough. Participants suggested something more visual that can be registered easily in the periphery and be understood at a glance such as a 'speed-o-meter' or 'power-bar'.

5.3 Study 2: Handedness using the Technique

The dominant hand has been established as more proficient in fine motor movements [11, 58, 268]. Typically in bimanual actions, the dominant hand is specialised for planning and controlling limb trajectory [329], with the non-dominant hand having advantages in force compensation and preserving posture [26]. Hand laterality is at its greatest in bimanual actions that require fine motor control, such as writing, with the non-dominant hand playing a supportive role [144]. Various studies have also shown that during symmetric bimanual actions with unilateral fatigue, task performance can be maintained with the non-fatigued hand compensating for the fatigued hand regardless of handedness [116, 160]. In Virtual Environments, as in the physical world, the hands are used both symmetrically and asymmetrically depending on the interaction mechanisms. For controller-based interaction, the most common interaction method using physical objects, both hands are designed to be tethered to the controllers holding them continuously while in the Virtual Environment. This is a limiting factor for prolonged and/or precise interactions. During actions, there is no opportunity for one hand to physically support the other leading to simultaneous fatigue of both hands.

Considering this there are two distinct advantages to the cube interaction technique. Firstly, the *configuration cube* is designed to be operated on a physical surface providing spatial stability during the interaction. This means that a user can manipulate the gain or scale of an object to a desired value with one hand, then switch tasks with the same hand while the *configuration cube* retains its state. This allows the user to switch between asymmetric and symmetric bimanual actions on the fly. The second key advantage is that the required precision of the hand using the *manipulation cube* can be altered on demand by adjusting the C/D gain. This is useful during complex manipulation tasks that require precision, but also useful for distributing the physical labour of both hands while maintaining the handedness arrangement during an interaction. For example, for prolonged interaction the hand using the *manipulation cube* will fatigue over time and precision will inevitably degrade. This can be mitigated by manipulating the gain factor and altering the required precision of the fatiguing hand.

With both these physical and virtual factors to aid with the distribution of labour, we theorise that the handedness arrangement should not impact task performance using the cube technique. For example, if the less precise non-dominant hand is using the *manipulation cube* the required amount of physical movement and precision can be controlled and compensated for by the dominant hand using the *configuration cube*. This is the case for translation and scaling tasks but not rotational tasks as there is a 1-to-1 mapping of the *manipulation cube's* orientation to the virtual object. It is also unclear if users will prefer a particular handedness arrangement or if user strategy will change due to the hands having naturally asymmetric roles in bimanual actions. To gain a better understanding of the interaction between the cube technique and handedness, we conducted a task-based behavioural study to answer these two sub-questions:

- Does handedness using the Cube Technique affect task performance?
- Do Users have a preferred handedness arrangement for the Cube Technique?

Similar to the previous study in Section 5.2, we compare handedness arrangements across a number of *small* and *large* virtual object manipulation tasks in AR. As in Section 5.2, the study involved users completing a series of virtual object 'docking' tasks [72, 104, 211, 264] where a virtual object is aligned to an identical target virtual object that differs in translation, rotation, and scale. Once again, we used the same low-poly Stanford bunny virtual model [356]. To measure task performance we record *task time*, total *hand movement* and *hand rotation*, and participant's *perceived workload* through NASA-TLX. We recorded participants' handedness preference, if there was one, and the reasoning behind their preference. Finally, we recorded both physical and virtual object positions for each trial to posthoc visualise participants' approach to aligning a virtual object over time using the cubes.

We recruited 10 participants (7 identified as male, 3 as female) with an age range of 22-40 (M = 30.5, SD = 6.1). 7 participants had used VR or AR HMDs occasionally, 1 participant reported weekly use, 1 reported daily use, and 1 participant had never used VR or AR HMDs before. 9 participants were right-handed and 1 participant was left-handed. The study followed a within-subject design with the handedness arrangement, translation size, scale size, and rotation size as factors. For the handedness arrangement, there were 2 levels: *Manipulation cube in the dominant hand* and*Manipulation cube in the non-dominant hand*. The same levels for the target translation size, rotation size, and scale size were used as in the pilot study: $smalltranslation$ = 10-99cm and $largetranslation$ = 100-250cm, $smallrotation$ = 0-50° and $largerotation$ = 70-180°, and $smallscale$ = 30-50cm and $largescale$ = 70-90cm. The chosen values and ranges for translation, scale, and rotation were motivated by the same factors in the pilot study — manipulations **within** and **beyond** comfortable reach — to encourage the use of the gain function during the study. The order of conditions was permuted using a counterbalanced Latin square design, with the target values being randomised with their specified ranges in the same manner as the pilot study. As a result, there were a total of $2 \times 2 \times 2 \times 2 = 16$ permutations with one set of repetitions for a total of 32 trials per participant ($32 \times 10 = 320$ total trials).

5.3.1 Apparatus, Tasks, and Procedure

The study was conducted with the same apparatus described in Section 5.2 using an HTC Vive Pro equipped with the Zed Mini video see-through AR camera, 60hz, resolution 2560 x 720, FOV 90° horizontal, 60° vertical, 100° diagonal max. The same study space was used — a large enclosed tracking space using the HTC Vive Lighthouses and a waist-height table for participants to use the interactive cubes on (shown in Figure 5.3). Participants' hands were fitted with two HTC Vive trackers. Two interactive cubes were used for this portion of the study and were calibrated using fiducial markers to the Vive tracking setup prior to the study, with the capacitive sensor calibrated while the cube is at rest on the table surface. The interactive cubes used in the study are as described in Chapter 4 Section 4.5. Data from the interactive cube's IMU and capacitive sensors and head and hands position/rotation were recorded at a sample rate of 60hz inside Unity. For this study, we updated the Heads-Up display and added a speed-o-meter style visualisation for the Gain factor (see Figure 5.3). We decided to use the **Absolute** gain design based on the results from our gain technique study (described in Section 5.2) and adjusted the minimum and maximum gain values to $G.MIN = 0.01$ and $G.MAX = 20.0$. We also constrained the handedness of participants for this study so that once a handedness arrangement was assigned, participants were not allowed to use both hands on either of the cubes — only the assigned hands.

Participants performed a series of alignment tasks in the same manner as the previous study. During a task, participants are shown two versions of the same virtual object and instructed to align the grey object (a control object which they manipulate with the cubes) over the blue transparent object (the target). Control objects were always automatically selected for the participant at the start of a trial. For this study, we applied automatic termination of a trial based on the pilot study results and the mean accuracy of users with the technique. A trial would be successfully completed and terminated by the system once the control object was aligned to the target object within $4cm$ positional difference, $6°$ rotational difference, and $3cm$ scale difference. The trial would be also be terminated if participants exceeded 90 seconds. As such participants were instructed to align the control object as close to the target as possible, with a secondary objective to be as quick as possible. We included the same visual cues to aid with alignment and placement which were projected next to the target displaying the distance, scale difference, and angular difference between the target and the control object (see Figure 5.3). The visual cues start to turn green once the control object is within a distance of *20cm*, an angular difference of $20°$, and a scale difference of *20cm*. The visual cues were the most saturated when at the minimum alignment values for position, rotation, and scale.

Prior to the study, participants completed a basic demographics form followed by a 10-minute induction to the study, AR system, and cube interaction technique. Participants would perform 2 sessions of trials, 16 trials per session, with 1 session for each handedness arrangement. Before starting a session of trials, participants were given as much time to practice the technique and tasks as they needed to become accustomed to the technique. Participants were free to take breaks between sessions and individual trials if necessary. Upon finishing a session participants were asked to complete the provided NASA-TLX form for that handedness arrangement. After both sessions participants completed a preference questionnaire to share their thoughts on their preferred handedness arrangement if they had one.

5.3.2 Study 2 Results

The measures for this study were task time, hand movement, hand rotation, and perceived workload and usability. We did not record alignment accuracy as trials were automatically terminated by the system once the control object was within $4cm$ positional difference, $3cm$ scale difference, and 6° rotational difference. The analysis was performed with a four-way repeated measures ANOVA ($\alpha = 0.05$) with Handedness (Manipulation cube in dominant vs. non-dominant hand), Translation Size, Rotation Size, and Scale Size as independent variables. QQ plots were used to validate the normality of the data. If the assumption of Sphereicity was violated as shown with Mauchly's test, Greenhouse-Geisser corrected values were used for analysis. In the case of significant results, Bonferroni-corrected posthoc tests were used. Effect sizes are reported as partial eta squared (η_p^2). RAW NASA-TLX scores were analysed using Friedman tests and, in the case of significant results, Bonferroni-corrected Wilcoxon signed-rank tests were used for posthoc analysis. Temporal filtering was used to remove outlier trials which we define as trials that exceeded the 90-second cap for completion. In total, 71 trials were discarded (22.19%) and after outlier filtering only 12 of the 160 cells had missing values for the four-way repeated measures ANOVA. These missing values were replaced with the maximum value over all participants for each measurement: task time, hand movement, and hand rotation. Any additional significant outliers identified in SPSS were winsorized. We also asked participants which handedness arrangement they preferred and to provide reasoning for their chosen ranking.

5.3.2.1 Task Time

To understand how handedness affects participants' ability to align a control object to a target using the cube interaction technique, we measured the task time for participants to complete an alignment. Task time is measured from the start of the trial to the moment when participants align the control object to the target within $4cm$ positional difference, 6.0° rotational difference, and $3cm$ scale difference. The results found no significant 4-way, 3-way, or 2-way interactions between any of the conditions when measuring task time. Likewise, there were no significant differences in task time between the 2 different handedness configurations ($F(1, 9) = 1.391$, $p > .05$, $\eta_p^2 = .134$). The manipulation cube in the **dominant hand** had a lower average task time ($M = 53.026$, $SD = 3.024$) as opposed to in the **non-dominant hand** ($M = 58.035$, $SD = 3.698$). However there was a significant main effect of translation size ($F(1, 9) = 52.220$, $p < .001$, $\eta_p^2 = .853$), with *small* translations were significantly faster ($M = 49.786$, $SD = 2.897$) than *large* translations ($M = 61.275$, $SD = 2.583$). This result corroborates our finding from the gain technique study with translation size significantly impacting task time.

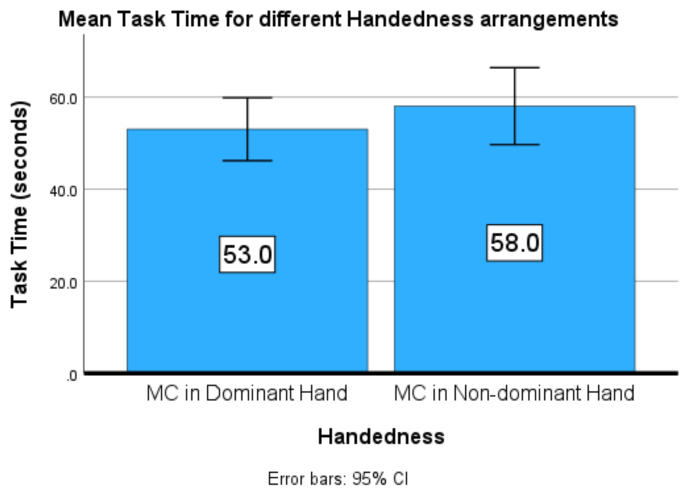

Figure 5.10: The mean task time for each handedness arrangement. $n = 10$

5.3.2.2 Movement and Workload

To see if different handedness arrangements influenced the total physical movement and user workload we measured the total physical hand movement and rotation of both hands using the hand-mounted Vive trackers and recorded participant's NASA-TLX scores. Total hand movement and total hand rotation are defined in the same way as the previous study (Section 5.2) — total hand movement is the total distance the hands moved during a trial in meters and total hand rotation is the tallied rotational movement of both hands in degrees. For total hand movement, the results found no significant 4-way, 3-way, or 2-way interactions between any of the conditions and, likewise, there were no significant main effects between the 2 different handedness arrangements ($F(1,9) = .725$, $p > .05$, $\eta_p^2 = .075$). The manipulation cube in the **non-dominant hand** had the highest average hand movement ($M = 2.091m$, $SD = .149$) versus the **dominant hand** ($M = 1.940m$, $SD = .113$). However, there was a significant main effect for different translation sizes ($F(1,9) = 14.811$, $p < .05$, $\eta_p^2 = .622$) and rotation sizes ($F(1,9) = 6.069$, $p < .05$, $\eta_p^2 = .403$). *Large translations* involved significantly more hand movement ($M = 2.180m$, $SD = .103$) than *small translations* ($M = 1.851m$, $SD = .111$) and the same was true for *large rotations* ($M = 2.105m$, $SD = .098$) versus *small rotations* ($M = 2.105m$, $SD = .111$).

For total hand rotation, the results found no significant 4-way, 3-way, or 2-way interactions between any of the conditions involving handedness and there were no significant main effects between different handedness arrangements ($F(1,9) = 1.973$, $p > .05$, $\eta_p^2 = .180$). However, there was a significant main effect for different rotation sizes ($F(1,9) = 16.780$, $p < .05$, $\eta_p^2 = .651$) with large rotations involving significantly

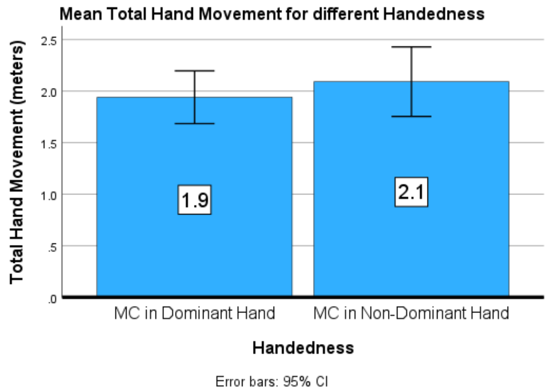

(a) The mean total hand movement for each handedness arrangement in meters.

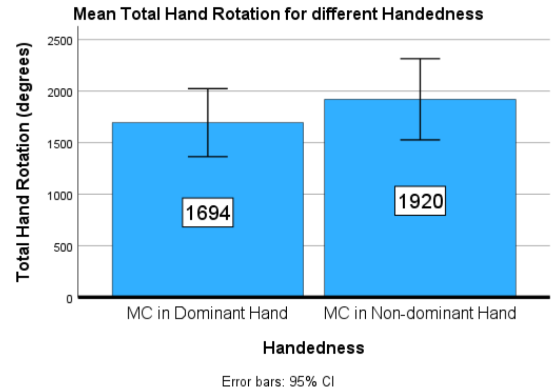

(b) The mean total hand rotation for each handedness arrangement in degrees.

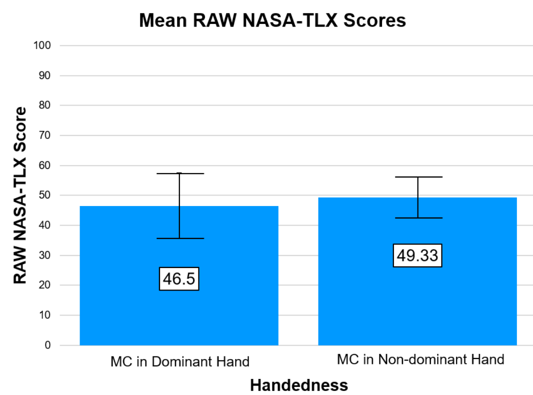

(c) The mean RAW NASA-TLX workload value for each handedness arrangement.

Figure 5.11: Mean movement and workload for each handedness arrangement. $n = 10$

more hand rotation ($M = 1663.598°$, $SD = 147.260$) than small rotations ($M = 1950.472°$, $SD = 138.960$).

Friedman tests using the RAW NASA-TLX workload metric revealed no significant differences between the different handedness arrangements for overall workload ($\chi^2(1, N = 10) = .500$, $p > .05$) although the manipulation cube in the dominant hand received a lower average workload score ($M = 46.500$, $SD = 17.576$) than the non-dominant arrangement ($M = 49.333$, $SD = 11.061$). Additionally, there were no significant differences for Mental Demand ($\chi^2(1, N = 10) = .500$, $p > .05$), Physical Demand ($\chi^2(1, N = 10) = .143$, $p > .05$), Temporal Demand ($\chi^2(1, N = 10) = .500$, $p > .05$), Performance ($\chi^2(1, N = 10) = .111$, $p > .05$), Effort ($\chi^2(1, N = 10) = .500$, $p > .05$), or Frustration ($\chi^2(1, N = 10) = .000$, $p > .05$).

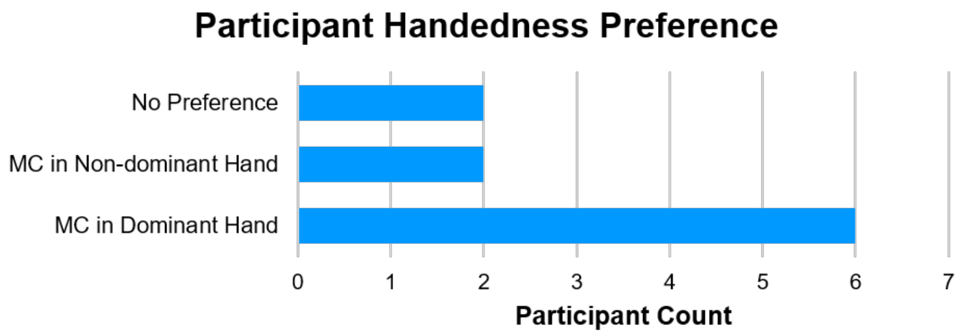

Figure 5.12: Total user preference for each handedness arrangement. $n = 10$

5.3.2.3 User preference and feedback

Figure 5.12 shows how many participants preferred each handedness arrangement. The majority of participants (6) preferred the *manipulation cube* in the dominant hand, 2 participants preferred the opposite arrangement, and 2 participants had no preference. Participants that voted for the *manipulation cube* in the **dominant hand** generally preferred it due to having finer motor control and less fatigue. P1 noted that the "*fine motorics*" of the dominant hand were better suited to "*fine-tuning the position*" of the control object, with P10 supporting this—"*It felt easier to maneuver the rabbit in my dominant hand*". Both P4 and P6 noted that using their **dominant hand** for manipulation felt "*intuitive*". P8 mentioned they felt "*more successful and quicker overall*" with their dominant hand, but also commented that they "*felt comfortable with both hands after a few trials*". P1 also mentioned not having a "*strong preference*". P7 suggested that their "*non-dominant hand fatigued a lot quicker*" than their dominant hand when using the *manipulation cube*. P7 also said that, despite their preference, they think they were "*more performant with the non-dominant hand*" using the *manipulation cube*, but speculated that this could be a "*learning effect*".

Participants (P3 and P9) that preferred the opposite arrangement, the *manipulation cube* in the **non-dominant** hand, decided this because they felt the scaling and gain control required the dominant hand due to them being novel or new interactions. P3 commented that the "*rotating and moving of the object was not as important*" while the "*gain regulation and scaling required more attention*" as it was a "*new interaction that required more conscious effort*". P9 supported this by saying the *manipulation* cube in the "*non-dominant hand was mentally less stressful*". Participants who had no preference, P2 and P5, still mentioned that their **dominant** hand had "*slightly better control*" but could not decide whether this would be better assigned to the **manipulation cube** or the **configuration cube**. Participants P6 and P7 noted that specifically "*scaling was easier with the dominant hand*" using the **configuration**

cube. P9 also shared this sentiment but attributed it to having a different mental model of how the scaling should work — "*my brain was thinking moving the cube right would make it [the object] bigger*".

Outside of preference, participants had an opportunity to share thoughts about the cube interaction technique as a whole. P1 and P4 both stated that the cube form factor "*influenced*" their usage of the technique. P1 specifically discussed how the dimensions of the cube would lead them "*to align the cube to the table as it felt the nicest in my hand*" while scaling. This would lead to unintended gain increases which were too large for the intended scaling. Interestingly, P1 said they were "*less aware of this when controlling the gain with their dominant hand*". P4 explains how they struggled to "*grasp at times during the task*" the mapping of the cube to the "*much more complex shape*" they were manipulating. P2 states that the most "*difficult aspect of the tasks was rotation*" alignment and P3 suggested that the "*gain should influence the rotation*". P5 suggests that handedness would not be an issue outside the experiment setup due to one cube being spatially stable on a physical surface—"*it makes it easy for whichever hand is holding the manipulation cube up and therefore it doesn't quite matter which hand is used*". Finally, P10 notes that "*the further it [the target] was, the harder I found the orientation in particular*", "*It felt much easier to accomplish the task when the rabbit was closer to me*".

5.3.3 Study 2 Discussion

Reflecting on the questions we posed for this study, our hypothesis was that handedness would not impact task performance using the cube technique which we measured as a product of task time, movement, and workload. From the results, we can see there were no significant differences between handedness arrangements in terms of task time, user hand movement, and self-reported NASA-TLX scores for virtual object alignment tasks. While the dominant hand has been established as more proficient at fine motor control in related work [11, 58, 268], our results demonstrate that for 3D manipulations using the cube technique, there was not a significant performance difference. This suggests that users may account for the reduced motor control in the non-dominant hand when using the manipulation cube by utilising the configuration cube to adjust the C/D gain to maintain performance during alignment.

As demonstrated by the similar NASA-TLX scores, explicit gain control and *algorithmic mapping* [221] can be applied to asymmetric bimanual interactions using cubic tools with the goal to more evenly distribute workload between the hands during different AR activities. This principle could also be applied to other form factors and may yield similar benefits, something which we explore in the next section. Although no significant differences in workload were found based on NASA-TLX scores, it remains unknown if both handedness arrangements experience equal levels of fatigue and if explicit C/D gain control has similar benefits for both handedness arrangements when using the cube technique. To further elucidate whether handedness performance can be equalised by giving the user control over C/D gain, future work could measure fatigue in the hand using the manipulation cube by using the consumed endurance model [164]. Additionally, if explicit gain control can maintain performance for users with less precise motor skills (i.e. when using the non-dominant hand), it could be a way to compensate for user fatigue over prolonged interactions.

Interestingly, we found that while users could not change the gain for virtual object rotation, there were no significant differences in task time between the handedness arrangements for different rotation conditions. This indicates that the cube's *re-graspability* can be utilized equally by both handedness arrangements. However, we observed that hand movement was significantly greater for large rotations compared to small rotations, suggesting that while the cube's *re-graspability* equalizes the time needed to align objects, it does not result in reduced hand movement.

Considering user preference, participants mostly preferred the **dominant hand** controlling the manipulation cube for virtual object alignments which could be a result of more precise motor control in the dominant hand [11, 58, 268]. However, multiple participants with this preference commented that this was not a strong preference with some describing how they became comfortable with both hands after a few trials. The participants that preferred the opposite arrangement reasoned that this was due to the gain and scale control being '*unfamiliar*' interactions that required more '*attention*'. This could be a product of user familiarity with VR as the participants that preferred this orientation used VR/AR occasionally or never. Likewise, this difference in preference could also be caused by different user approaches to the alignment tasks. For example, some users may choose to focus their efforts on regulating the gain more with the configuration cube to reduce the work of the manipulation cube.

Outside of handedness, our study uncovered several insights into the cube interaction technique itself. Participants reported that the cube's form factor influenced their use of the technique, with some suggesting that the cube's *clear dimensions* [235] led them to align the configuration cube more with the physical surface paying less attention to the gain factor value. Another participant suggested that the mapping between the cubic tool and virtual object was difficult to understand due to their very different apparent form factors, and another described how *re-grasping* the manipulation cube was difficult. Despite explicit gain control enabling users to manipulate virtual objects with less movement, our results showed a significant difference in task time and movement for different translation sizes, which was also observed in the previous study.

Finally, there were several limitations with our study opening up avenues to explore the cube interaction technique further. For example, while we gained an understanding of expected user accuracy using the cube technique in Study 1, employing automatic task termination in this study reduces our understanding of the upper limit of user accuracy with different handedness arrangements. We partially infer this upper limit from participant task time, it is possible that higher accuracy can be achieved with one-handedness arrangement over the other if task time was unlimited. As mentioned previously, future work could look at measuring the fatigue of different handedness arrangements using the consumed endurance model, but also future work could consider more longitudinal AR activities in context. While it would have not been possible in our study setup, another avenue for future research could be to further explore the benefits of explicit gain control by comparing techniques with and without this function. This would further elucidate the advantages of explicit gain control, regardless of the interaction or applied AR task.

5.4 Study 3: Technique Comparative Evaluation

After having explored different facets of the cube interaction technique in depth, we now consider how the cube technique compares to state-of-the-art manipulation techniques for AR. We conducted a comparative evaluation using the cube technique and two interaction techniques designed for the HTC Vive Controllers (described in Figure 5.13 and 5.14). The first controller technique was operated *unimanually* and the other *bimanually*. To meaningfully compare the techniques, both controller techniques were designed to have the same capabilities as the cube technique in terms of virtual object manipulation and gain control. This was to ensure that users were still able to complete the manipulation tasks with each technique. We also refrained from implementing the controller techniques with widgets or menus and relied purely on the input mechanisms provided on the physical devices. The controller techniques are described in Section 5.4.1.

Specifically, this study looks at how the cube technique performs across *small* and *large* virtual object manipulations when compared to the two controller techniques. We leveraged the same 'docking' tasks described in the previous two studies (Section 5.2 and 5.3) and used the same Stanford bunny [356] virtual model for all tasks. To measure task performance we record *task time*, total *hand movement* and *hand rotation*, and participant's *perceived workload* through NASA-TLX. We also asked participants to rank the three interaction techniques based on their preferences and explain their reasoning. Finally, we recorded both physical and virtual object positions during a trial to posthoc visualise participants' approach to aligning a virtual object using the three different techniques.

We recruited 12 participants (8 identified as male, 4 as female) with an age of 25-44 (mean: 31.58, SD: 5.62). 7 participants had used VR and AR head-mounted displays occasionally, 3 participants had never used VR/AR HMDs, 1 participant reported weekly use, and 1 participant reported daily use. 11 participants were right-handed and 1 participant was ambidextrous. The study followed a within-subject design with the interaction technique, translation size, scale size, and rotation size as factors. For the interaction technique, there were 3 levels (*cube technique, unimanual controller,* and*bimanual controller*). The same levels for the target translation size, rotation size, and scale size were used as in the gain technique and handedness study: $smalltranslation$ = 10-99cm and $largetranslation$ = 100-250cm, $smallrotation$ = 0-50° and $largerotation$ = 70-180°, and $smallscale$ = 30-50cm and $largescale$ = 70-90cm. The chosen values and ranges for translation, scale, and rotation were motivated by the same factors — manipulations **within** and **beyond** comfortable reach [73, 174] — to encourage the use of the gain function during the study for all 3 interaction techniques. The order of conditions was permuted using a counterbalanced Latin square design, with the target values being randomised with their specified ranges in the same manner as the gain technique and handedness study. As a result, there were a total of $3 \times 2 \times 2 \times 2 = 24$ permutations with one set of repetitions for a total of 48 trials per participant ($48 \times 12 = 576$ total trials).

5.4.1 Apparatus and Techniques

The study was conducted with the same apparatus as the two prior studies. We used an HTC Vive Pro equipped with the Zed Mini video see-through AR camera, 60hz, resolution 2560 x 720, FOV 90° horizontal, 60° vertical, 100° diagonal max. The same study space was used — a large enclosed tracking space using the HTC Vive Lighthouses and a waist-height table for participants to use the interactive cubes on. Participants remained seated at the table throughout the study and their hands were fitted with two HTC Vive trackers. Based on the results of the handedness study (described in Section 5.3) we allowed participants to decide their handedness arrangement for each technique which they maintained throughout the trials. All data from the interactive cubes, controllers, hand trackers, and AR headset were recorded at a sample rate of 60hz inside Unity. In order to understand the advantages and disadvantages of the cube form factor, e.g. using a cube with a physical surface to adjust the gain and scale factors, and manipulating virtual objects unimanually and bimanually, we compared our technique to two controller-based methods. The different technique conditions are as follows:

5.4.1.1 Cube Technique

Two interactive cubes were used for this study and were calibrated using fiducial markers to the Vive tracking setup prior to the study, with the capacitive sensor calibrated while the cube is at rest on the table surface. The interactive cubes used in the study are described in Section 5.1 and the operation of the technique is described in

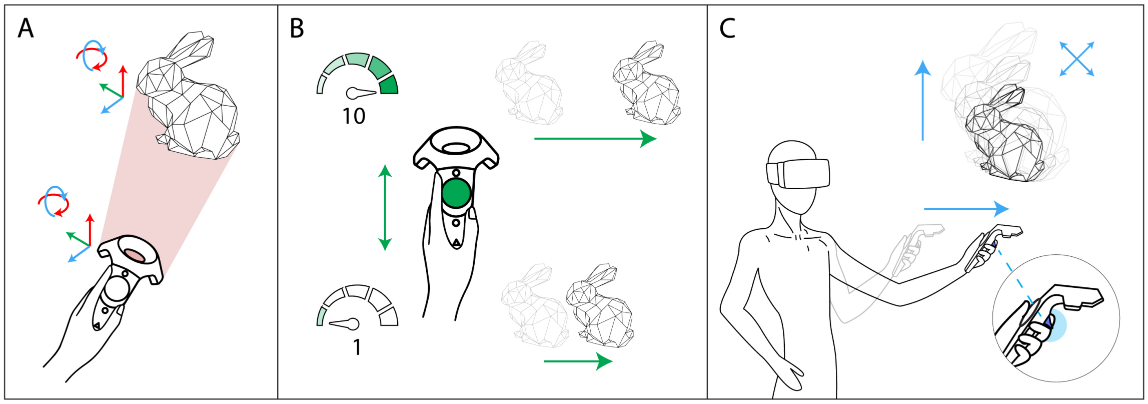

Figure 5.13: The unimanual controller technique using an HTC Vive VR controller. **A)** The controller operates as a spatial proxy for a virtual object in translation and rotation. **B)** Pressing or holding *Up* on the D-pad increases the gain and *Down* decreases the gain. This influences the translation and scale but the rotation retains a 1-to-1 mapping. **C)** The scale is increased by holding the *Trigger*, which freezes the object in place, and moving the controller away from the user (stretching) [428]. Scale is decreased in the same manner but moving the controller towards the user (squeezing) [428].

Section 5.1.1. We decided to use the **Absolute** gain design based on the results from our gain technique study (described in Section 5.2). The min and max gain values were set to $G.MIN = 0.01$ and $G.MAX = 20.0$. We also constrained the handedness of participants based on our results from the handedness study (Section 5.3) but participants could decide their handedness before the study — i.e. the dominant hand assigned to the *manipulation cube*. To visualise the gain factor, the same speed-o-meter visual was used on the Heads-Up display (See Figure 5.3).

5.4.1.2 Unimanual Controller

This technique, shown in Figure 5.13, uses a single HTC Vive controller as a proxy for the virtual object with the object mimicking the position and rotation of the physical controller. The controller is designed to primarily be operated in the dominant hand, but the non-dominant hand can be used to support during a manipulation. Similar to the cube technique, there is a gain factor that is visualised on the HUD as a speed-o-meter and this can be increased by pressing/holding Up on the D-pad and decreased with Down on the D-pad. To scale a virtual object, the user must press and hold the trigger to enter a 'scaling mode' where the virtual object freezes in place. The user then moves the controller away from themselves to enlarge the object and towards themselves to shrink the object — all while holding the trigger. This

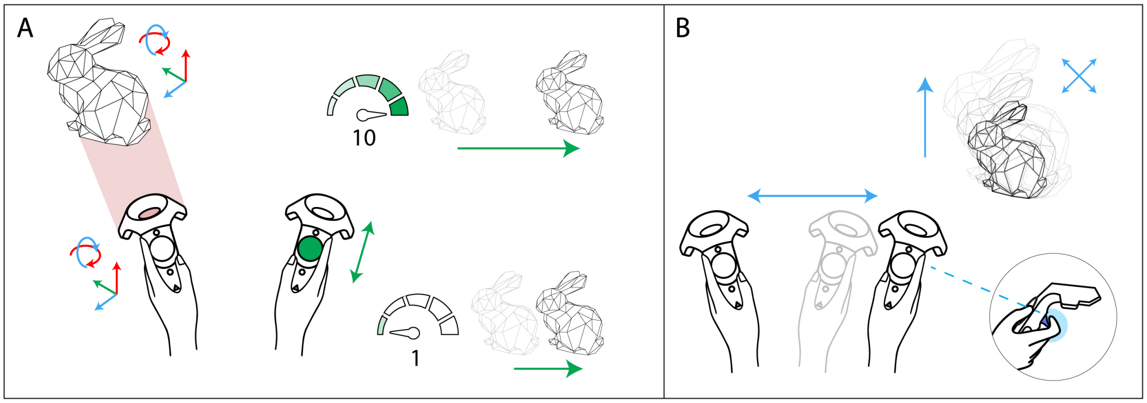

Figure 5.14: The bimanual controller technique using two HTC Vive VR controllers. **A)** One controller operates as a spatial proxy for a virtual object in translation and rotation. The other controller provides system controls for increasing and decreasing the gain by pressing *Up/Down* on the D-pad. **B)** The scale is increased by holding the *Trigger* on the controller that controls the gain, which first freezes the object in place. Then moving the controllers away from each other increases the object scale (stretching) [428]. Scale is decreased in the same manner but moving the controller towards each other(squeezing) [428].

scaling method leverages a similar 'stretching' and 'squeezing' metaphor implemented in the cube technique but instead uses the position of the controller relative to the user. Both the position and scale of the virtual object are influenced by the gain factor but the rotation retains a 1-to-1 mapping with the controller. As mentioned the participant must operate the controller in their dominant hand and they can use their non-dominant hand to support, but during the study, they were instructed to not switch the controller to a different hand.

5.4.1.3 Bimanual Controller

This technique, shown in Figure 5.14, uses two HTC Vive controllers operated in separate hands. This technique operates similarly to the cube technique, with the controller in the dominant hand as the *manipulation* controller and the controller in the non-dominant hand as the *configuration* controller. The virtual object mimics the position and rotation of the *manipulation* controller. As with the other two techniques, there is a gain factor that is visualised on the HUD as a speed-o-meter and this is manipulated with the *configuration* controller. The gain is increased by pressing/holding Up on the D-pad and decreased with Down on the D-pad. Scaling of a virtual object is done by pressing and holding the trigger of the *configuration* controller to enter a 'scaling mode' where the virtual object freezes in place. The object is then scaled by moving the *configuration* controller away from the other controller to enlarge the object, and towards to shrink the object — while holding the trigger. This scaling method leverages a similar 'stretching' and 'squeezing' metaphor implemented in the cube technique. Both the position and scale of the virtual object are influenced by the gain factor but the rotation retains a 1-to-1 mapping with the *manipulation* controller. During the study, participants were instructed to not switch either controller to a different hand. The *manipulation* controller must be operated in their dominant hand and the *configuration* controller in their non-dominant hand.

5.4.2 Tasks and Procedure

Participants performed a series of alignment tasks in the same manner as the gain technique and handedness study. During a task, participants are shown two versions of the same virtual object and instructed to align the grey object (a control object which they manipulate with the cubes) over the blue transparent object (the target). Control objects were always automatically selected for the participant at the start of a trial. We applied automatic termination of a trial based on the gain technique study results and the mean accuracy of users with the cube technique. A trial would be successfully completed and terminated by the system once the control object was aligned to the target object within $4cm$ positional difference, $6°$ rotational difference, and $3cm$ scale difference. The trial would also be terminated if participants exceeded 90 seconds. As such participants were instructed to align the control object as close to the target as possible, with a secondary objective to be as quick as possible. We included the same visual cues to aid with alignment and placement which were projected next to the target displaying the distance, scale difference, and angular difference of the target and the control object (see Figure 5.3). The visual cues start to turn green once the control object is within a distance of *20cm*, an angular difference of 20°, and a scale difference of *20cm*. The visual cues were the most saturated when at the minimum alignment values for position, rotation, and scale.

Prior to the study, participants completed a basic demographics form followed by a 10-minute induction to the study and AR system. Participants would perform 3 sessions of trials, 16 trials per session, with 1 session for each interaction technique. Before starting a session of trials, participants were given as much time to practice an interaction technique with some example trials to become accustomed. Participants were free to take breaks between sessions and individual trials if necessary. Upon finishing a session participants were asked to complete the provided NASA-TLX form for that particular interaction technique. After all three sessions, participants ranked the three interaction techniques and shared their thoughts in a questionnaire on their chosen ranking.

5.4.3 Study 3 Results

The measures for the comparative evaluation were task time, hand movement, hand rotation, and perceived workload and usability. As in the handedness study we did not record alignment accuracy as trials were automatically terminated by the system. Termination was triggered once the control object was within $4cm$ positional difference, $3cm$ scale difference, and $6°$ rotational difference. The analysis was performed with a four-way repeated measures ANOVA ($\alpha = 0.05$) with Interaction Technique (Cubes, Unimanual Controller, and Bimanual Controller), Translation Size, Rotation Size, and Scale Size as independent variables. QQ plots were used to validate the normality of the data. If the assumption of Sphereicity was violated as shown with Mauchly's test, Greenhouse-Geisser corrected values were used for analysis. In the case of significant results, Bonferroni-corrected posthoc tests were used. Effect sizes are reported as partial eta squared (η_p^2). RAW NASA-TLX scores were analysed using Friedman tests and, in the case of significant results, Bonferroni-corrected Wilcoxon signed-rank tests were used for posthoc analysis. As in the other studies, temporal filtering was used to remove outlier trials which we define as trials that exceeded the 90-second cap for completion. In total, 90 trials were discarded (15.63%) and after outlier filtering only 16 of the 288 cells had missing values for the four-way repeated measures ANOVA. These missing values were replaced with the maximum value over all participants for each measurement: task time, hand movement, and hand rotation. Any additional significant outliers identified in SPSS were winsorized. We also asked participants to rank the interaction techniques based on their preference and provide reasoning for their chosen ranking.

5.4.3.1 Task Time

As in the two previous studies, we recorded task time to compare the cube interaction technique efficiency against the two controller techniques when aligning the control object to a target. Task time is measured from the start of the trial to the moment when participants align the control object to the target within $4cm$ positional difference, $6.0°$ rotational difference, and $3cm$ scale difference. When analysing the normality of the results, we found a positive skew in the distribution. We accounted for this by transforming the data using the square root of the aggregated means purely for statistical analysis. The actual recorded means are reported otherwise.

The results found no significant 4-way interaction, however, there was a significant 3-way interaction for $InteractionTechnique \times TranslationSize \times ScaleSize$ ($F(2, 22) = 6.759$, $p < .05$, $\eta_p^2 = .381$) and $InteractionTechnique \times RotationSize \times ScaleSize$ ($F(2, 22) = 6.193$, $p < .05$, $\eta_p^2 = .360$). For $InteractionTechnique \times TranslationSize \times ScaleSize$, the pairwise comparison showed that for *small translations* and *large scales* the **cube** technique was significantly slower ($M =$

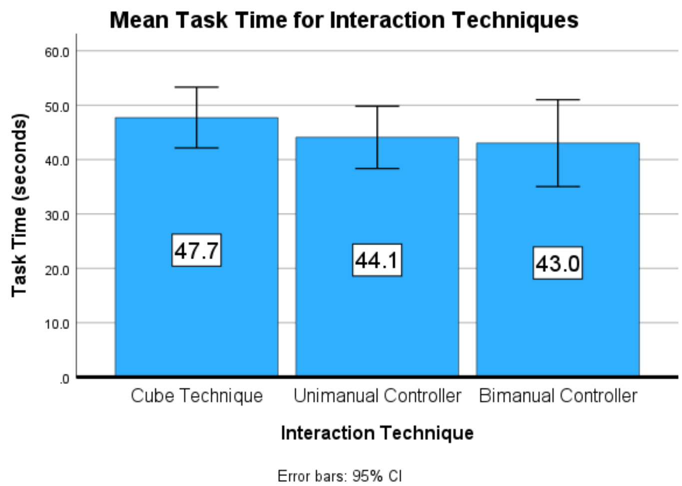

Figure 5.15: The mean task time of each interaction technique. $n = 12$

42.991, $SD = 2.593$) than the **unimanual controller** technique ($M = 34.718$, $SD = 3.075$). **All techniques** were significantly slower for *large translations* versus *small translations* for conditions with both *small scales* and *large scales*. The **unimanual controller** was the only technique to be significantly faster for *small scales* ($M = 46.008$, $SD = 3.385$) versus *large scales* ($M = 56.983$, $SD = 3.243$) with a *large translation*.

For $InteractionTechnique \times RotationSize \times ScaleSize$, the pairwise comparison showed that for *small rotations* with *large scales* the **cube technique** was significantly slower ($M = 53.200$, $SD = 3.570$) than the **bimanual controller** technique ($M = 40.362$, $SD = 3.487$). For the **bimanual controller** technique *large rotations* with a *small scale* ($M = 40.601$, $SD = 3.682$) were significantly faster than with a *large scale* ($M = 48.259$, $SD = 3.482$). Additionally for the same technique, *large scales* with *small rotations* ($M = 40.362$, $SD = 3.487$) were significantly faster than *large scales* with *large rotations* ($M = 48.259$, $SD = 3.482$). Surprisingly, the inverse was true for the **cube technique** with *small rotations* ($M = 53.200$, $SD = 3.570$) being significantly slower than *large rotations* ($M = 44.022$, $SD = 2.861$) with *large scales*. The results found no significant 2-way interactions involving different techniques however, as observed in the prior studies, **translation size** was a significant main effect ($F(1, 11) = 91.518$, $p < .05$, $\eta_p^2 = .893$), with *large translation* ($M = 51.703$, $SD = 3.074$) taking significantly longer than *small translations* ($M = 38.198$, $SD = 2.205$). **Interaction Technique** was not a significant main effect on task time ($F(2, 22) = 1.889$, $p > .05$, $\eta_p^2 = .147$).

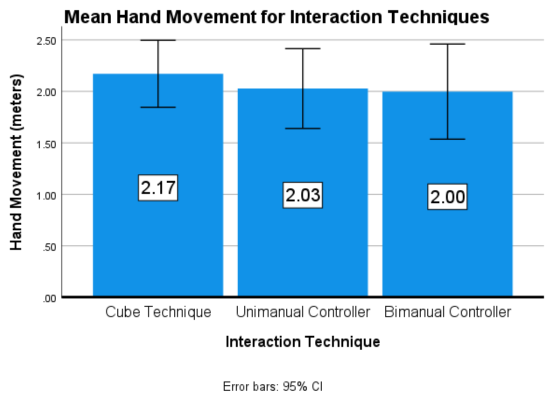

(a) The mean total hand movement for each interaction technique in meters.

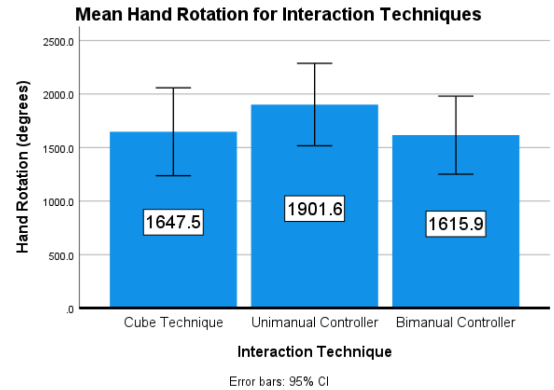

(b) The mean total hand rotation for each interaction technique in degrees.

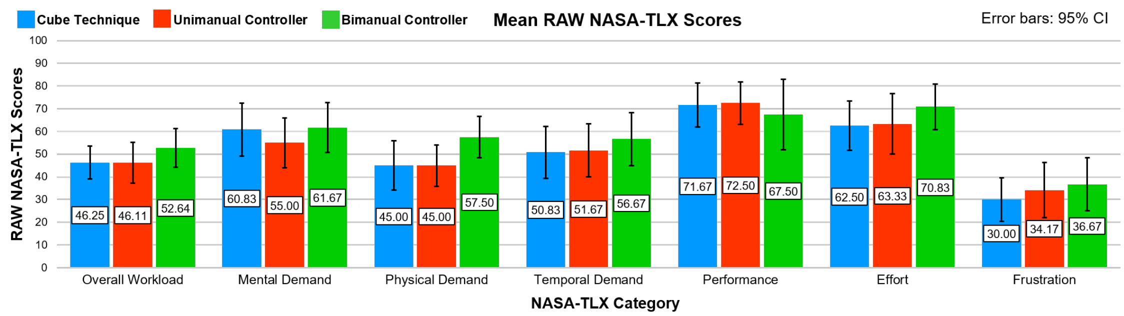

(c) The mean RAW NASA-TLX scores for overall workload and each workload type per interaction technique.

Figure 5.16: The mean user movement and workload for each interaction technique. $n = 12$

5.4.3.2 Movement and Workload

To compare the relative physical movement and user workload between the three techniques we measured the participant's total physical hand movement and rotation of both hands using the hand-mounted Vive trackers. We also recorded participants' responses to the NASA-TLX questionnaire for each technique. Total hand movement and total hand rotation are defined in the same way as the two prior studies — total hand movement is the total distance the hands moved during a trial in meters and total hand rotation is the tallied rotational movement of both hands in degrees. When analysing the normality for both hand movement and hand rotation, we found both sets of results had a positive skew in distribution. The hand movement data was transformed using the square root of aggregated means and the hand rotation used the logarithm of aggregated means, both of which resulted in a more normal distribution across conditions. Again, the transformation is purely for statistical analysis with the actual recorded means reported.

For total hand movement the results found no significant 4-way interaction but, as in the task time analysis, there was a significant 3-way interaction for $InteractionTechnique \times TranslationSize \times ScaleSize$ ($F(2, 22) = 12.390$, $p < .05$, $\eta_p^2 = .530$). Looking at the pairwise comparison showed that for *small translations* with *large scales* the **cube** technique required significantly more hand movement ($M = 1.369m$, $SD = .058$) than the **unimanual controller** technique ($M = 1.168m$, $SD = .047$). The **cube** technique required significantly more hand movement for *small scales* with *large translations* ($M = 1.494m$, $SD = .064$) versus *small translations* ($M = 1.279$, $SD = .041$), which was also the case for *large scales* with *large translations* ($M = 1.564m$, $SD = .068$) versus *small translations* ($M = 1.369m$, $SD = .058$). For the **unimanual controller** technique, this was only true for *large scales* with *small translations* ($M = 1.168m$, $SD = .047$) versus *large translations* ($M = 1.562m$, $SD = .102$). For the **bimanual controller** technique, the inverse was true with *small scales* with *small translations* requiring significantly less hand movement ($M = 1.171m$, $SD = .061$) versus *large translations* ($M = 1.424m$, $SD = .076$). The results also found a significant 2-way interaction for $InteractionTechnique \times RotationSize$ ($F(2, 22) = 4.036$, $p < .05$, $\eta_p^2 = .268$). However, the only significant result from pairwise comparison showed that the **bimanual controller** technique had significantly more hand movement for *large rotations* ($M = 1.397m$, $SD = .068$) versus *small rotations* ($M = 1.295m$, $SD = .059$). As in the previous studies and with task time, there was a significant main effect involving *translation size* ($F(1, 11) = 31.033$, $p < .05$, $\eta_p^2 = .738$), but none for *interaction technique* ($F(2, 22) = .986$, $p > .05$, $\eta_p^2 = .082$).

For total hand rotation, no significant 4-way interaction was found but the same significant 3-way interaction was found for hand movement for $InteractionTechnique \times TranslationSize \times ScaleSize$ ($F(2, 22) = 4.476$, $p < .05$, $\eta_p^2 = .289$). This time the pairwise comparison showed no significant simple main effects between the interaction techniques. Otherwise, similar main effects were observed for hand rotation as in hand movement. The **cube** technique required significantly less hand rotation for *small translations* versus *large translations* for both *small scales* (small translation - $M = 1367.67°$, $SD = 174.35$, large translation - $M = 1735.00°$, $SD = 213.90$) and *large scales* (small translation - $M = 1520.27°$, $SD = 216.71$, large translation - $M = 1967.09°$, $SD = 250.87$). For the **unimanual controller** technique, there was significantly less hand rotation in *large scales* with *small translations* ($M = 1495.44°$, $SD = 197.85$) versus *large translations* ($M = 2444.47$, $SD = 317.85$). The **bimanual controller** technique had significantly less hand rotation in *small scales* with *small translations* ($M = 1252.19°$, $SD = 142.69$) versus *large translations* ($M = 1778.57°$, $SD = 209.84$). The results found no significant 2-way interactions but there were two significant main effects involving *translation size* ($F(1, 11) = 24.011$, $p < .05$, $\eta_p^2 = .686$) and *rotation size* ($F(1, 11) = 5.508$, $p < .05$, $\eta_p^2 = .334$), but no significant effect of *interaction technique* ($F(2, 22) = 1.109$, $p > .05$, $\eta_p^2 = .092$).

Friedman tests using the RAW NASA-TLX workload metric revealed no significant differences between the different interaction techniques for the overall workload ($\chi^2(2, N = 12) = 1.167$, $p > .05$), see Figure 5.16c. Additionally, there were no significant differences for Mental Demand ($\chi^2(2, N = 12) = .531$, $p > .05$), Physical Demand ($\chi^2(2, N = 12) = 4.043$, $p > .05$), Temporal Demand ($\chi^2(2, N = 12) = 1.050$, $p > .05$), Performance ($\chi^2(2, N = 12) = .974$, $p > .05$), Effort ($\chi^2(2, N = 12) = 2.513$, $p > .05$), or Frustration ($\chi^2(2, N = 12) = .195$, $p > .05$).

5.4.3.3 User preference and feedback

Based on participant's preferred ranking (shown in Figure 5.17), 5 participants most preferred the **cube technique**, 5 participants for the **bimanual controller**, and 2 participants for the **unimanual controller**. Considering the comments of participants that ranked the cube technique the highest, 3 participants mentioned the benefits of using the cube on a physical surface. P2 mentioned that "*there is a gorilla arm problem*" with the controller techniques and that the controllers were "*quite clunky*". P2 emphasises "*it is nice that the gain cube is on the table, controlling the effort of the other arm*". P1 supports this— "*It didn't feel natural to put one controller down if I wanted to use both hands for rotation. However, in the cube technique, it felt fine to leave one cube on the table and help my dominant hand*". Additionally, 4 participants mentioned that they preferred the cube specifically for rotating and orientating the virtual object. P1, P6, P8, and P12 all mentioned that the cubes were "*best for rotating*" with P1 also noting that the "*controllers were troublesome for rotating*". However, P1, P2, and P8 stated that they preferred the controllers for the button-based scaling and gain control but this was not enough to rank the controller techniques higher.

Participants that preferred the **bimanual controller** the most expressed two reasons for their ranking— They preferred the separated and button-based input for scaling and manipulating the gain, and the controller techniques were more responsive. P3, P4, and P10 described how the buttons and triggers provided "*finer control*" over the gain and scaling "*than the cubes*". P3 specifically said the "*unimanual controller was the most awkward*" with too many controls "*all in one hand*". P10 also referenced the better separation of gain and scale in the **bimanual controller**—"*I had more control over each separate function*". Other participants, P3 and P5, suggested that the higher latency of the cube technique resulted in ranking the bimanual controller technique higher—"*the cubes are good but lack a fast response*" (P5) and "*there is more delay*" (P3). However 4 participants mentioned that they preferred the cube

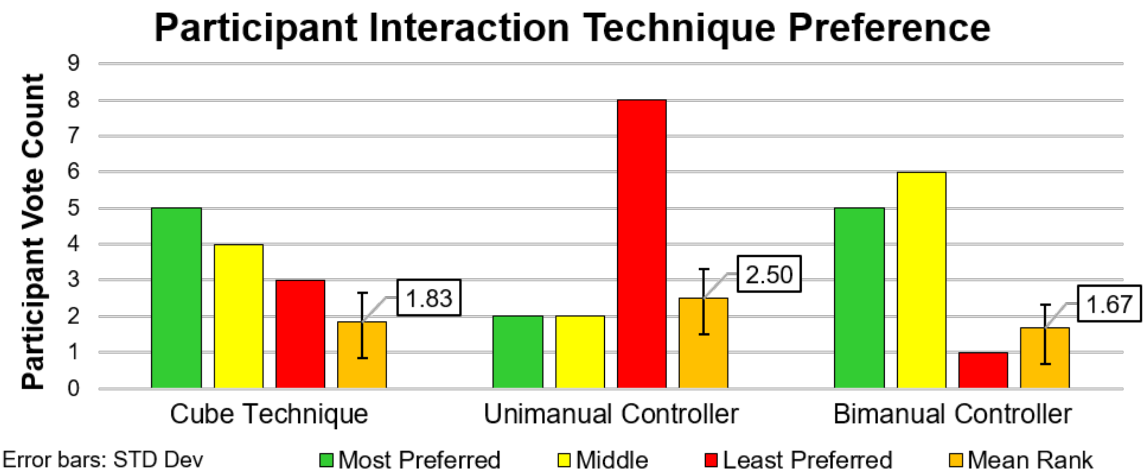

Figure 5.17: User preference for each interaction technique from most preferred to least preferred, as well as the mean ranking. $n = 12$

technique specifically for rotating and moving the virtual objects—"*cube interaction was the most comfortable for manipulating position*" (P4), "*cubes were a lot better for rotation since the shape moved much more easily in my hand*" (P9). Participants also criticised the controller form factor for rotations—"*I needed to use my other hand to support the controller in positions where the weight was unevenly distributed*" (P9), "*controller is awkward to manipulate when rotating and moving*" (P3). The two participants that preferred the **unimanual controller** the most also mentioned that having the button-based control for the gain and scale was beneficial but also that having the one hand free to support the other was preferable—"*I could use one hand purely to support my other when doing micro gestures*" (P11). Once again, both participants mentioned that the **cube technique** was preferable for rotating and translating—"*The cubes felt slightly easier to handle*" (P11).

When providing additional comments, 8 out of 12 participants discussed combining techniques and altering the cube form factor to improve the cube interaction technique. 4 participants suggested combining the controllers and the cubes with the cube manipulating the control object and the controller used to refine gain or sometimes scale. P1 said "*Ideally, I would have a cube in my dominant hand and a button interaction (for the gain) in my non-dominant hand*" and P4 supported this—"*I would have like a controller/cube combo as physical buttons are easier for the fine adjustments*". P10 commented "*the cubes would have been my preferred method if there was a button to activate the gains control and deactivate the scale control like on the controllers*". P2 mentioned that the 'freezing' of the object by pulling the trigger in the controller methods was useful and would be useful in the cube technique

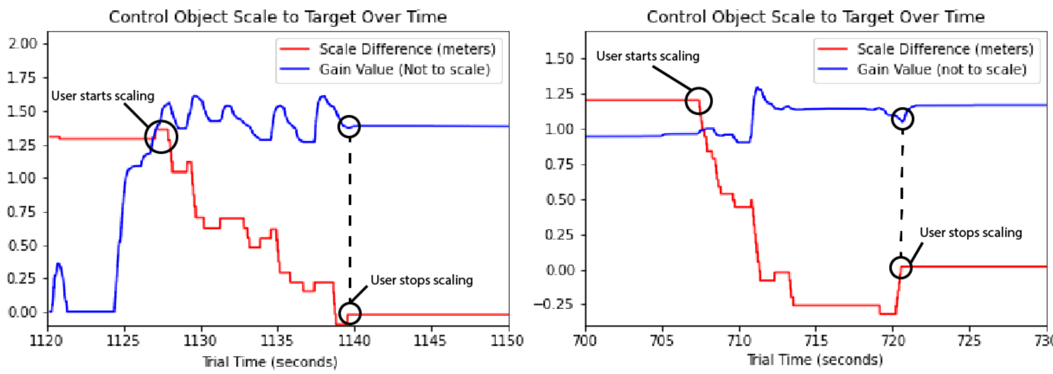

Figure 5.18: Two example trials of gain and scale functions interfering with one another when using the *cube technique*.

"*allowing you to move your arm if the gain is low*" more during scaling. Another 4 participants suggested altering the cube form factor or technique. P11 suggested that the graspability would be better if "*cubes were smaller*" and P12 said that the uniform nature of the cubes made it difficult for them "*to get a sense of orientation*" in relation to the control object. P12 followed this up with "*if it had something that I can take as a reference point*" instead of feeling out the rotation. P3 and P6 both suggested that the "*gain and scale*" should be "*decoupled*" more, with P3 suggesting a "*modal-switch*".

5.4.3.4 Trial Observations

Considering the performance measures analysed and the qualitative feedback from users, we highlight observations from exemplar trials. As observed in the gain technique study (section 5.2), scaling the control object and adjusting the gain using the cube technique would occasionally interfere with one another shown in Figure 5.18. Despite the refinement of the gain after the gain technique study, opting for the *absolute* design, the close coupling of scale and gain in the cube technique appears to induce unintentional interactions for some users in certain target conditions. Figure 5.18-right is an example of how unintentional gain changes can sometimes lead to overshooting the target.

After reviewing participants' qualitative feedback on their technique preference, we compared the cube technique with the controller techniques when performing particular manipulations shown in Figure 5.19. Particularly, participants discussed gain and scale control as something they preferred in the controller techniques and Figure 5.19-left shows an example of scaling the control object in an exemplar trial. The scaling in the cube technique was often slower and more incremental than the

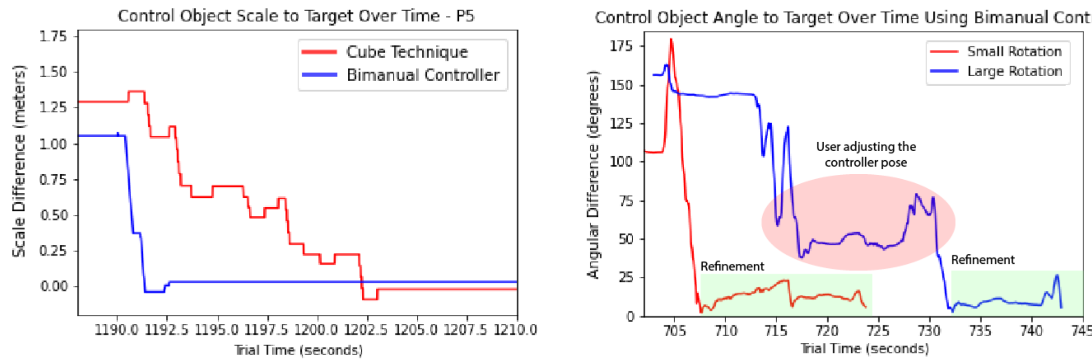

Figure 5.19: **Left** compares the *cube technique* and the *bimanual controller* scaling behaviour from two separate trials for P5. In the cube technique, scaling is more incremental over a longer duration whereas the bimanual controller is more direct over a shorter duration. **Right** compares the *bimanual controller* behaviour when aligning rotation to the target for small and large rotational differences. In *small rotations*, the controller can directly reach the target rotation with minimal re-grasping of the controller, followed by refinement. In *large rotations* the participant reaches the threshold of their wrist rotation and has to re-orientate the controller with one or both of their hands which is time-consuming.

controller techniques which were often quicker. However, a large proportion of the participants preferred the cube techniques for moving and rotating. In relation to this, Figure 5.19-right shows a small versus large rotation for the bimanual technique. During large rotations, participants often had to account for the non-uniform shape of the controller and had difficulty orientating the controller in hand to achieve the target rotation which would often be timely (more than 15 seconds in Figure 5.19-right). This was particularly prevalent for the bimanual technique as users could make use of their free hand in the unimanual technique to support the rotation.

Despite the controller techniques often being more direct in scaling than the cube technique, a very prevalent behaviour in all participants was overshooting the target when scaling (see Figure 5.20). This was the case for both controller techniques but was mostly observed in scaling manipulations, not translation or rotation. This was not observed as much when scaling with the cube technique. This could suggest that designing one gain factor to simultaneously influence the two types of manipulation, scale and translation, could be sub-optimal for the controller techniques. However, the correction/recovery from overshooting the target was typically direct and, as the lack of significant differences in task time between the techniques suggests, did not impact participant performance.

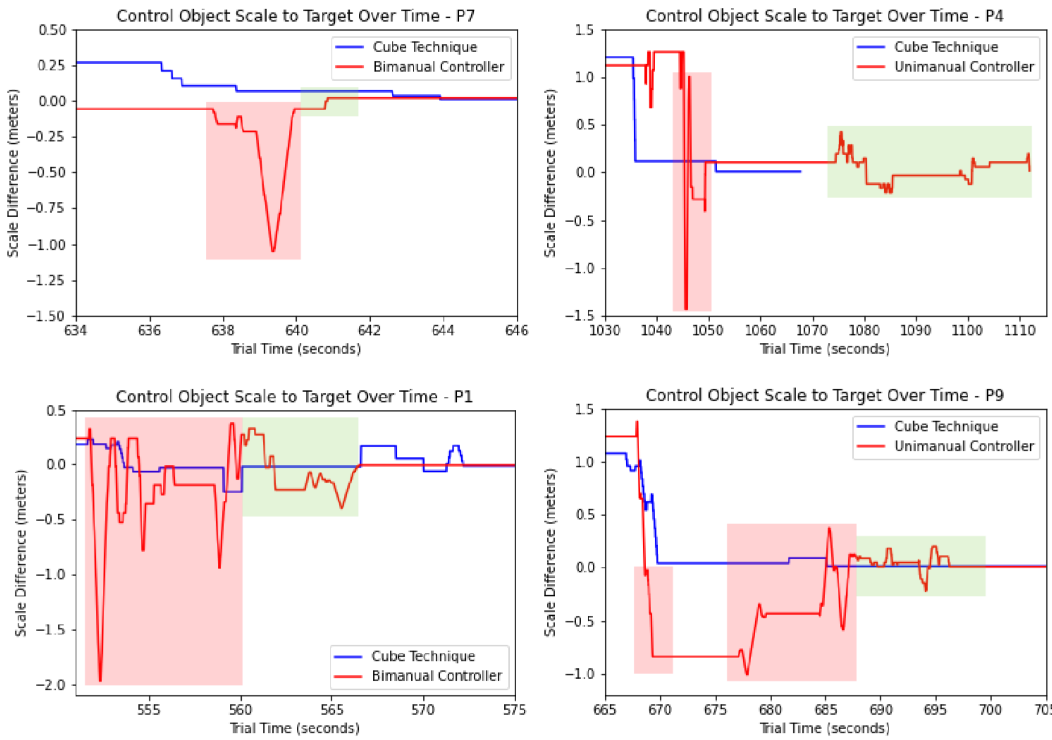

Figure 5.20: Several example trials from P1, P4, P7, and P9 demonstrating large overshoots (highlighted in red) followed by refinement (highlighted in green) when scaling the virtual object using the *bimanual* and *unimanual* controller techniques.

5.4.4 Study 3 Discussion

After thoroughly exploring cubic tools and the designed interaction technique in the two previous studies, let us directly address **RQ3** following the results of our comparative evaluation. Our findings show that participants performed similarly in virtual object alignment tasks using both the designed cube technique and the state-of-the-art hand-controller techniques. However, participant preference was divided, with some preferring the bimanual controller technique and others the cube technique. Further analysis reveals several points for discussion and potential conclusions.

Firstly, a number of performance metrics and observations suggest that the scale and gain functions were too closely coupled in the cube technique. Considering task time, the unimanual and bimanual controller techniques were significantly quicker on average in alignment conditions with larger scaling. For movement, this was also the case when comparing the cube technique to the unimanual controller also in alignment conditions with larger scaling. In our posthoc observations we found that for some participants the gain and scale were influential on each other, in a similar manner to what we found in Study 1. This is also supported by participant comments with many expressing that the gain and scaling control was preferred in the controller conditions, with more control over '*each separate function*', and some even explicitly stating gain and scale should be decoupled in the cube technique. For our specific technique and AR tasks, this could suggest that *spatial handles* are not as preferred or proficient for **system control** [221] versus discrete input afforded by buttons and triggers on controllers, but this would require further exploration and evaluation. On the other hand, the unimanual controller was also criticised for its lack of separation between the scale and gain, despite being assigned to separate input elements, but this was only mentioned by P3.

Secondly, beyond **system control**, most participants said they much preferred the cube technique as a proxy for the virtual object versus the controller techniques. Several participants noted that the cube was much easier for rotation, suggesting that '*re-graspability*' played a role. This aligns with the findings of Englmeir et al. [111] who found a sphere form factor was preferred for manipulation tasks over standard VR/AR controllers. Participants described the controller form factors as "clunky" for proxy-based interaction and noted that their weight was unevenly distributed, making manipulations more challenging. The controllers also did not afford placement on the physical surface provided in the study space. In the bimanual controller technique, even when one controller was not being used, participants still held it, which sometimes hindered the alignment task. To support this, our observations showed an example of a participant having to compensate for the difficult form factor of the controller during rotation. Although participants preferred the cube as a proxy, our quantitative measures did not show a significant difference in performance. As to our design, many participants felt comfortable moving and adjusting the configuration cube on the physical surface and also felt comfortable leaving it on the surface to control the manipulation cube with both hands, something which is afforded by cubic tools and physical surfaces.

Thirdly, it is possible that the closely coupled gain and scaling control in the cube technique is not the limiting factor, but rather a limitation of the technique's implementation. In our study, the cube technique used simple thresholds on acceleration, as determined by pilot testing and previous studies. However, this resulted in some unintentional interactions due to differences in user movements especially when controlling scale and gain. To improve the robustness of the technique, a machine-learning classifier could be applied to the fusion of sensor data, similar to our toolkit's surface gesture recognition. Some participants also mentioned experiencing more latency when rotating with the cube technique compared to the controller techniques, but despite this, they maintained a strong preference for the cubes for rotation due to its '*re-graspability*'.

Our observations of the controller techniques found that participants were often overshooting the target when scaling. Unlike the limitations of the controller form factor which is difficult to address to improve rotation, this overshooting is likely down to the design of the interaction technique. As mentioned prior, it could be that there would be improved performance if the scale and translation had separately controlled gain values. On the other hand, the technique scaling could be too sensitive for some users. Either way, overshooting the target scale did not appear to have a significant impact on task time or movement as participants were often quick to recover.

Fourthly, our study revealed that participants preferred the controller techniques for **system control**, specifically gain control, while preferring the cube technique for **3D manipulation** tasks like translating and rotating the virtual object. This suggests that a combined form factor, incorporating the best of both techniques, could potentially provide a more optimal experience. This idea was also supported by the study participants, with several expressing that they would have preferred a controller for gain control while using a cube for manipulation.

The study aimed to evaluate the performance of basic cubic tools that are virtually modified compared to state-of-the-art controller techniques. Future research could consider evaluating physically-modifiable tools with a combination of form factors, ranging from simple shapes like cubes and spheres to more complex form factors such as those conceptualised in Chapter 3 (Section 3.5). Furthermore, our study aimed to investigate the use of PM-Tools that can be reconfigured *between* interactions. However, future research could delve into exploring physical modification *during* interaction as an input mechanism, as previously explored in Chapter 3. Additionally, while Chapter 4 presented several ways tools can be virtually modified through various interaction metaphors, the methods for virtual modification discussed in Chapter 3, such as using handles or containers, have not yet been empirically investigated.

Lastly, further research is needed to understand the impact of explicit gain control on 3D interaction in AR, as mentioned previously in Section 5.3.3. It remains unclear if different tools would benefit differently from access to gain control. For instance, cubic tools may benefit more or less compared to hand controllers, which may be better suited to using raycast techniques for manipulation. This would shed light on whether virtual capabilities are more or less influenced by the form factor of a tool and if it is capable of physical modification, which would be consistent with our understanding of perceived affordance [178, 282, 284].

To summarise, there are several noteworthy advantages and disadvantages between the techniques. The cubic tool technique performed very similarly to the controller-based techniques according to our performance and self-reported measures. It is also important to note that while cubic tools have the capacity to support controller-style interaction techniques for 3D manipulation, they provide unique interactions not usually afforded to traditional AR controllers such as combining with the physical environment to provide a spatial *handle* and virtual *containers*. This highlights the need for further exploration and refinement of current AR controllers with our results showing that even basic tool form factors, such as a cube, can provide competitive results and in some cases outperform.

5.4.5 Summary

In this chapter, we delved deeper into a specific interaction technique for precise 3D manipulation, implemented using the physically-modifiable cubic tools developed in Chapter 4. We answered **RQ3** by first empirically exploring the design and behavioural phenomena around the technique, with respect to gain design and handedness, and then directly comparing the cubic tools against AR/VR hand-controllers inspired by the current state-of-the-art techniques. Our three lab studies yielded several conclusions.

Firstly, for virtually modifiable tools, there are a number of different ways to design an interaction that may result in better performance. In our first study, we looked at three different methods for implementing gain control using a cubic spatial *handle* all utilising a 'dial' metaphor on a physical surface. While we expected some differences in performance between the different gain designs, our results found only minor differences between the techniques. However, this study was crucial to our understanding of the average efficiency and accuracy of the cubic tools for precise 3D manipulation. User behaviors, such as synchronous and asynchronous alignment and gain 'clutching', were also identified.

Our second study evaluated the impact of handedness on the use of bimanual, asymmetric cube tools. Results showed that different handedness arrangements performed similarly due to explicit gain control and precision in manipulating a cube bimanually. These findings were not only important for informing the design of our cube technique for the comparative evaluation, but they also pose interesting questions for future research. For example, applying gain control to other tools and interactions, and using gain control to mitigate and compensate for physical fatigue.

Thirdly, we explicitly addressed **RQ3** by performing a comparative evaluation to understand how the cubic tool interactions compare with state-of-the-art AR techniques such as hand controllers. We can conclude that the cubic tools performed similarly and in some cases outperformed the hand-controller techniques due to their proficiency in '*re-graspability*'. We also highlighted the preference for discrete input mechanisms for system control, over using spatial handles, and the opportunities for combining tool form factors in future work.

The next Chapter reflects on the work presented in this thesis, revisits the research questions, and discusses limitations and opportunities for future work.

Chapter 6

Discussion

In the preceding chapters, we explored the use of physical tools for 3D interaction in AR through the different concepts described by Napier [275]: *Tool-Making*, *Tool-Modifying*, and *Tool-Using* for AR. In this chapter, we revisit our proposition of physical tools as the salient form of interaction for the future of AR and reflect on the research questions we posed in Chapter 1. We then discuss the outcomes and limitations of the work and studies throughout the thesis. Finally, we look at the opportunities for future work to expand knowledge of physical tools for labour in virtual spaces and provide our final reflections on current interaction devices and the envisioned future of Tools in Augmented Reality.

6.1 Research Questions

The field of HCI has a rich history of developing tools for interacting with machines [31]. This thesis has built on this prior work by exploring physical *Tool-Making*, *Tool-Modifying*, and *Tool-Using* in AR by positing and investigating the research questions described in the introduction. In Chapter 1, we also asserted that (1) interaction with virtual objects in AR would be frequent and essential, and (2) that the creation, modification, and adoption of physical tools was vital in addressing current interaction challenges. The examination of related work in 3D interaction [55, 221], haptics [94, 305], and tangibles [46, 189] in Chapter 2 support these assertions and showcases the benefits of physical devices to mediate interaction with virtual content in AR and VR [17, 111, 117, 273, 434]. The work and findings in this thesis also reinforce the benefits of using physical tools for interaction in AR, as demonstrated by our conceptual examples (Section 3.5), demonstrative applications (Section 4.4 and 4.6), and designed interaction techniques (Section 5.4.1). Before revisiting our research questions, let's first reflect on these unique benefits that arose in prior Chapters.

Physical AR tools can offer a tangible and intuitive means to handle complex interactions, including 3D manipulation and modeling of virtual objects, demonstrated by the studies exploring cube-based interactions in Chapter 5. Beyond static physical tools, our exploration of modifiable tools in Chapter 3 shows how expressive yet precise interactions can be facilitated through an interaction device's capacity to *augment* or *reconfigure*, expanding the capabilities of simple form factors such as a *cube*. Our design space of cubes in AR in particular demonstrates that the physical attributes of a device shape the design of its virtual interactions and the metaphors employed. Our resulting prototypes and applications using cubes demonstrate how AR tools can be designed to leverage the physical environment, such as surfaces, to increase precision when interacting with virtual objects, serve as a pedestal for creating system controls, and provide scaffolding for organising virtual workspaces. Now let's consider our Research Questions outlined initially.

6.1.1 RQ1: *Tool-Making* - How do we ideate tools for AR that are both physically and virtually modifiable?

In Chapter 3, we presented a framework for describing and classifying physically-modifiable **objects** (PM-Objects) which we use as a vehicle to explore physically-modifiable **tools** (PM-Tools) in AR using a combination of pre-existing methodologies such as guessability and research through design. Our findings revealed a general agreement on mapping *reconfigurable* and *augmenting* PM-Object characteristics to AR interactions, which we operationalised in PM-Tool conceptual examples, but there was less consensus on interaction design and limited consideration of the physical environment. The proposed framework provides a starting point for designing PM-Tools in AR. Insights from applying the framework and accompanying studies provide avenues to explore tool modification further and highlight the importance of intentionally incorporating the physical environment in the design process.

6.1.2 RQ2: How should physically and virtually modifiable tools be created and operated for AR? (*Tool-Modifying*)

Based on the insights from answering the prior research question, we explored a specific form factor for AR interaction, cubes, intentionally incorporating physical surfaces from the outset. For directly answering **RQ2**, we produced a toolkit (TangibleTouch) enabling the fabrication of instrumented and modifiable cubic tools. The cubic tools can detect a user's surface gestures and are capable of being tracked in 3D, which facilitates an array of potential approaches to operating the tools in AR. We illustrate the broad design space and feasibility of *TangibleTouch*, and PM-Tools generally, through the different interaction metaphors and demonstrative applications discussed in Chapter 4.

6.1.3 RQ3: How do newly designed tools compare to existing interaction techniques and methods for AR? (*Tool-Using*)

By using the *TangibleTouch* toolkit as a platform we were able to directly address **RQ3** by comparing the performance and preference of cubic tools against state-of-the-art hand-controllers. In Chapter 5 we focussed on a specific interaction technique using cubic tools, which was 3D object manipulation. We then provided an in-depth exploration and evaluation of the technique design itself, the behavioural phenomena of handedness using the technique, and a comparative evaluation of VR/AR controllers. We gained insights into the distribution of labour in asymmetric bimanual interactions by providing system control and direct manipulation that allowed users to adopt different strategies when striving for precision. We also found that cubic tools performed similarly to our designed hand-controller techniques for 3D manipulation, but that crucially the techniques excelled in very different areas; the cube was widely preferred for direct manipulation, with the controllers preferred for system control. These insights support the findings of others on AR/VR interaction using basic form factors [111] and also open up new and interesting avenues for future work when designing physical tools for AR.

6.2 Outcomes of the Thesis

We derived several outcomes from the work presented in the thesis which serve as interesting topics of discussion but also considerations for designing physical and modifiable tools for AR in the future.

6.2.1 Tool-Modification in AR is Nuanced

In addressing the research questions presented in the thesis we found that the concept of physically-modifiable tools is more nuanced than first expected. In Chapter 2 we grounded PM-Tools in the context of work in shape-changing interfaces [321], tangible artifacts [338], and modifiable toolkits and proxies [111, 117, 226]. In Chapter 3 we explored PM-Tools from the perspective of leveraging their modifiable properties for interaction, in other words using their affordances as direct mechanisms for input, yielding a number of agreed-upon mappings to AR interaction which we embodied in our concepts of PM-Tools. However, reflecting on related work around modifiable toolkits [117] and our prototype presented in Chapter 4, physical-modification can also be viewed as a process of changing the purpose and function of a tool from one thing to another.

Initially, we approached Tool-modifying as not just changing how a tool appears, physical modification, but also how it is applied, virtual modification. We considered the process of changing an AR tools function to be predominantly an issue of mapping to virtual content — virtual modification. However, as shown in our work with the *TangibleTouch* toolkit and our cubic tool prototypes in Chapters 4 and 5, sometimes a tool has to be physically modified for new capabilities to be realised at which point it can be virtually modified.

Furthermore, PM-Tools can be designed not only for different contexts and applications but also for different types of users in mind. Chapter 3 explored physical modification at the behest of users *during an interaction* with virtual objects to enable expressive yet precise manipulations. Whereas Chapter 4 explored PM-Tools with modification *between interactions*, such as changing the faces of the same cubic tool to enable different surface gestures mapped to different functions in an application, which can be leveraged by users and designers alike. For example, users can use modification between interactions to repurpose the same form factors for different applications, whereas designers can rapidly test different tools and interaction designs. Reflecting on the concepts of *reconfiguring* and *augmenting* PM-Tools introduced in Chapter 3, we could speculate that these characteristics may be more or less suited to *during* and *between* modification for interaction, which we expand on in Section 6.3 as an avenue for future work.

6.2.2 The Physical Environment is central to AR & Tools

One of the central arguments presented in the introduction was that the physical environment is often not actively considered in AR/VR interaction design and can even be an obstacle at times. This thesis maintains that leveraging the physical environment yields opportunities for new and improved interactions in AR. We also assert that direct interaction with virtual objects would be frequent and necessary for the future of AR, but even if this is not the case, the physical environment remains a central and distinguishing aspect of AR, influencing concepts such as pervasive and nomadic AR [142, 143, 213] determining where and how information is overlaid.

Our exploration of cubic tools revealed that incorporating them with physical surfaces enhances interaction through added sensory cues in addition to the passive haptics innate to tangible objects. This facilitated interaction techniques that met and in some cases exceeded the precision of current state-of-the-art AR techniques and by providing surface support, mid-air arm fatigue is reduced. Additionally, the combination of tools and physical surfaces, as shown in Chapter 4 with cubic tools, allows for an ad-hoc configuration of system controls that mimic familiar physical interfaces such as dials, sliders, and buttons.

The physical environment and surfaces also serve as a scaffold for organizing a virtual workspace, labour, and objects. However, while physical tools are required to leverage the environment as an organisational structure for defining virtual areas and organising heterogeneous information and virtual objects, it is important to note that not all physical tools can leverage the physical environment in the same capacity. As demonstrated in our work on PM-Tool design in Chapter 3, despite including physical surfaces as part of the design process it was not always utilised in the design of gestures and interactions. Additionally, our comparative evaluation in Chapter 5 evidenced how VR/AR hand controllers conflicted with the physical environment and surfaces, with users not even placing down controllers when it would be beneficial to do so. The utilisation of the physical environment is dependent on the form factor and affordance of a tool, whether it can be physically modified, the design of interactions, and the tool's sensing capabilities.

The physical environment can be leveraged for interaction in different ways, as shown by the considerable related work on 'opportunistic controls' [156, 157, 159, 403] and the work in this thesis on AR tools. We approached this relationship through the lens of interactive artifacts for AR activities such as 3D manipulation and modeling. Our findings highlight the advantages of tools that enable operation on physical surfaces, offering flexibility in interaction design and supporting different types of user labor, such as direct interaction with virtual objects, indirect surface-supported labor, or a combination of the two. Furthermore, our work serves as an exemplar of the active consideration that should be given to aspects of the physical environment during the process of tool-making and modification. For instance, we highlight the use of a cubic form factor, specifically chosen for its affinity towards physical surfaces [235, 339], as a deliberate design decision.

Future work can delve into the design of tools tailored to specific physical environments, exploring further methods to integrate the physical environment or existing instruments of physical labor. Additionally, there is scope to expand upon opportunistic controls, offering virtual capabilities that are context-specific. A more comprehensive exploration of these possibilities is detailed in Section 6.3, where we discuss these ideas in greater depth.

6.2.3 Reflecting on Making & Using Tools in AR

Considering the framework and methodologies we employed for tool creation in Chapter 3, we explored the mapping of physically-modifiable objects (PM-Objects) affordance to AR interactions. Hornecker describes the difficulty in mapping affordance to digital interaction [178] and as such in our work we probed affordance mapping through empirical studies not to derive absolute rules for designing tools in AR, but rather as an exploratory approach to their development. Chapter 3 evidences how probing affordance mapping with users has the potential to distill useful characteristics and properties from existing physical objects that can then be combined into new form factors for AR tools.

The findings across the thesis show that *augmenting* PM-Tools can act as a mediator to support natural hand gestures as a guide for more expressive and precise input by providing crucial haptic sensory cues. Whereas *reconfigurable* PM-Tools can provide efficient and granular input for system control, in a similar capacity to hand-controllers, that can be used in conjunction with direct gesture. However, as discussed previously, gesture and interaction design lacked agreement amongst users. This could imply that there is some collective understanding and consensus around the applicability of physical characteristics of tools, but that the design and operation of the tool itself and its relationship to virtual content is disputed between individuals. This also may be due to the limitations of the guessability methodology we employed, something which we expand on further in the next section.

Reflecting on our examination of physical AR tools, our findings in Chapter 5 show that even simple forms, like cubes, can compete with, outperform, and be preferred over standard techniques like hand controllers. Specifically, cubes were found to be particularly favored for directly rotating virtual objects, which was similar to the findings of Englmeier et al. [111] and performed equally to designed controller methods. Our synthesis of related work in Chapter 2 highlights that tangible time-multiplexed interfaces are more widespread in AR and VR than space-multiplexed interfaces [46] such as standard AR/VR hand-controllers that perform various functions across applications. The design of our cubic tools has shown that space-multiplexed interfaces offer clear advantages when combined with physical surfaces to divide virtual workspaces and provide simultaneous control over different virtual functions, such as one cube as a virtual object proxy and another to manipulate the control-display gain. Our cubic tools have the potential to be both space and time-multiplexed interfaces, which can be further studied regarding tool multiplicity and form factor diversity to construct larger interaction device ecosystems, i.e. a toolbox of physical AR tools. Our work is just the beginning, and we delve into future prospects in the following section.

6.3 Limitations & Future Work

This research has shown the potential of physically-modifiable tools for AR through various examples, techniques, and applications. While it showcases the numerous opportunities to explore physical tools in virtual environments, there are still areas that this initial investigation has not delved into and which deserve further attention.

This section outlines those limitations and directions for future research.

6.3.1 Tool-Making in AR

Our exploration of physical tools for AR worked in the context of head-mounted display AR which has considerable generality to other mixed reality mediums such as VR, mobile AR, smart glasses, and projection AR. However, there are different factors to consider for tool making between different AR mediums, for example, in mobile AR the hands are usually busy holding the display. In this case, our approach to tool-making could be used but instead for physical tools integrated into the mobile devices themselves such as the work of Visschedijk et al. [399].

Furthermore, our work explored a subset of canonical AR tasks, 3D interaction, modelling, and system control. Future work could explore physical tools for other basic AR tasks such as information browsing, but also in domain-specific applications in domestic and industrial settings. Often times the context of work limits the resources available in a virtual environment and in some cases the focus may be on physical labour with the virtual environment in the periphery. These stringent factors necessitate and guide the development of more bespoke physical tools or the appropriation of existing objects and tools for virtual interaction.

Additionally, our development of physical tools for AR was guided from the perspective of a single user, but often times virtual environments facilitate collocated and remote collaboration between many users. While this work has touched on egocentric and exocentric interactions in AR and discussed the potential of tools facilitating work between multiple users, this requires further exploration and should be explicitly considered in the design process of AR tools.

Examining the physically-modifiable objects and tools framework and studies presented in Chapter 3, there are a number of limitations that require further exploration. Our studies did not result in many gestures and interactions with PM-Tools that involved the utilisation of the physical environment. This could be a result of the participant's cultural understanding and preconceptions of the example PM-Objects we used to probe PM-Tools in AR, thus limiting the types of interactions designed. Additionally, the consensus set of user gestures had a low consensus when compared to other guessability studies. It could be that other methodologies should be employed to explore user-defined tangible gestures beyond current elicitation approaches, which have known and widely discussed limitations [385].

6.3.2 Tool-Modifying in AR

In designing modifiable tools, our conceptual focus was split between physical-modification *during* and *between* interactions. Our work specifically centered on the fabrication, prototyping, demonstration, and evaluation of PM-Tools that allow for modification *between* interactions. To fully realise the potential of physical modification as an interaction mechanism, as demonstrated in our conceptual examples from Chapter 3, technical implementations and prototypes must be developed and evaluated.

Our toolkit, *TangibleTouch*, provides a tracking and sensing framework that can be adapted to any physical form factor. For instance, applying the same method of tracking the hands to 3D track any instrumented physical object, and utilising capacitive sensing to support user surface gestures. While there are technical challenges in prototyping new tools with different physical modification capabilities, future research should build on the modular approach used in *TangibleTouch* and our cubic tools to maintain the essence of accommodating both designers and users.

In addition to the cubic form factor, there are several avenues to explore different form factors, particularly in the context of combining them, as discussed in the previous chapter (Section 5.4.4). One promising direction for further exploration is the combination of cubes with sphere proxies, like the work of Englmeier et al. [111], which offers an immediate opportunity to delve deeper into the multiplexing of physical tools using primitive form factors.

Moreover, the incorporation of shape-change actuation mechanisms [320, 373] presents another avenue for investigating the effectiveness of other primitive shapes, such as planes, spheres, and prisms, for proxy-based manipulation of virtual objects. This exploration could also involve deploying multiple physical proxies at scale to prototype interactions, akin to the work conducted by Feick et al. [117]. By considering these possibilities, we can push the boundaries of physical AR tool design and refine our understanding of how different primitive form factors and shape-change capabilities can enhance interaction in virtual environments.

6.3.3 Tool-Using in AR

This thesis explores and evaluates a specific interaction technique, using the developed cubic tools, for precise 3D manipulation of virtual objects. However, as discussed in Chapter 4, there are many other potential interactions and AR activities that could be enabled by cubic and physical tools. Further research is needed to fully understand these other interaction designs through in-depth exploration, including comparative and formal evaluations with users to assess their performance and effectiveness. Our previous evaluation only compared the cubes to a specific form factor of hand-controllers, and a more comprehensive comparison with other techniques, such as mid-air gestures or alternative controller form factors (e.g. Valve Index [391], MagicLeap [247]), would provide valuable insights. While the modifiable tools from the TangibleTouch toolkit have been shown to be versatile, they have not been formally evaluated by designers or users. Future work could address this to further solidify the utility of the toolkit.

As mentioned prior, cubes afford being passed between users and, as emphasised in Activity Theory [53, 218], objects and tools impact the way we work collaboratively. As such the potential of cubic tools in multi-user collaboration, both in co-located and remote settings, should be evaluated. For example, comparing collaborative AR tasks using current AR interaction techniques to those using the cubic tools, or exploring the use of cubic tools in remote collaboration such as the 'portals' mentioned in Section 6.4. Also, throughout the thesis, we have mentioned that physical objects and tools can aid in mediating interaction between users with different display capabilities, but this would need to be formally explored in future work.

A clear outcome from the work presented in Chapter 5 is that combining different tool form factors should be more thoroughly explored for AR activities. Further research could investigate the potential of an AR toolbox; combining the concepts of proxies, handles, and containers in different tools; or expanding the exploration of cubic tool techniques beyond 2 or 3 cubes. The thesis draws inspiration from Beaudouin-Lafon's work on Instrumental Interaction [31] and the concept of 'meta instruments' which we apply to the cubic tools, but this is a preliminary discussion that requires further exploration. For example, future work could further investigate the authoring and modification of tools by other tools, and even examine context-awareness in physical tools, such as adapting the tool's form based on the user's location and activity.

6.4 Reflections on Tools, AR, & the Future

Throughout the preceding chapters, research has been conducted to explore physical and modifiable interaction devices aimed at enhancing the user experience in AR. However, it is important to note that the primary focus of this thesis is maintained on the concept of "physical tools." The body of work conducted in the development of this thesis has provided a more comprehensive and profound perspective of work in virtual environments, leading to a shift in emphasis from "physical devices" to "physical tools." As we reflect upon the insights garnered from our studies, it becomes crucial to distinguish between the current landscape of **AR interaction devices** and provide a definition of the envisioned future of **AR tools**. Below are the broader defining factors of future AR Tools that arose from this body of work.

6.4.1 Process of Creation

Interaction devices are essential for engaging with computers, whether through desktops or AR interfaces. The development of these devices typically follows a cyclical process involving designers creating them, users adopting them, and empirical research providing insights for new refinement, designs, or practices. This process is also applicable to modern commercial and consumer-based physical and digital tools present in our work and domestic spaces. However, our historical relationship with tools has not always followed this pattern. In the past, tool users were intimately involved in the design and creation process, acquiring knowledge and skills to modify their tools based on user experience or evolving needs. This allowed them to refine and enhance the tools over time.

In modern times, users have become increasingly detached from the process of tool-making and modifying. While this abstraction is often seen as convenient, it has resulted in a diminished sense of autonomy and intimacy in the work performed with these tools. Moreover, if users lack the ability to create and refine their own tools, it can present a barrier to adoption when the available tools do not meet their specific needs or expectations, particularly in terms of diversity and accessibility. Looking ahead to the future of AR tools, being a tool user to the fullest extent also means being an empowered tool maker and modifier, equipped with the necessary means to design, research, and refine one's own tools and practices.

While this perspective can be relevant to various contexts and environments beyond just AR, we emphasize its significance within the realm of AR. With its potential for pervasiveness and integration into established practices, AR holds boundless possibilities to enhance capabilities within specific contexts where generic tools would fall short. However, designing and developing specialized AR tools at scale can be challenging. Therefore, empowering users with the means to create and refine their own AR tools has the potential to distribute and democratize AR on a broader scale. By adopting this approach, future AR tools should embrace a transient nature, continuously flowing between usage, design, and evaluation. As a consequence, the traditional boundaries between users, designers, and researchers blur and diminish, fostering a collaborative and iterative ecosystem.

6.4.2 Mechansims for Refinement

Considering the anticipated proliferation and pervasiveness of AR technology in work and domestic environments, it is crucial to acknowledge that barriers arise for physical tool-making and modification in the same capacity as learning current interaction devices and techniques. It becomes essential to establish mechanisms that empower users of all backgrounds to reflect, evaluate, and enhance their AR tools seamlessly, thereby cultivating mastery within the virtual environment. This thesis proposes three approaches to facilitate this process for physical tools in AR.

Firstly, incorporating **shape-change** mechanisms in various forms offers the potential to imbue tools with organic or expressive characteristics when leveraged *during* interaction, and provide flexibility and broader utility if leveraged *between* interaction. Just as AR adapts to our physical surroundings, our tools should similarly possess adaptive qualities. Secondly, **modularity** in physical AR tools allows the construction of more complex devices using commonly understood components, allowing for tools to expand sensing or haptic capabilities, creating an inherently multi-modal form of interaction. Lastly, recognizing the importance of **multiplexing** in our physical tool use, where tools often have intricate relationships with one another (e.g., a hammer and chisel), it is essential to support such relationships in virtual environments. By doing so, we can unlock congruous and compounding capabilities that enhance the overall functionality of the tools within virtual environments.

This perspective immediately raises questions regarding the fabrication and evaluation of such tools, which currently present challenges in terms of practicality and accessibility. Further insights are required to address these challenges effectively. However, the increasing adoption of domestic 3D printing and fabrication techniques suggests that open-source hardware and software toolkits offer a promising and immediate pathway for the dissemination of physical AR tools, serving as a stopgap while broader approaches are developed.

Expanding beyond the scope of this thesis, an essential aspect of this new perspective on future physical AR tools is fostering an attitude towards AR and virtual environments that goes beyond individualized portals to disconnected, bespoke manufactured worlds. Instead, it emphasizes the importance of collective and communal spaces that promote equal and open access to tools and the knowledge surrounding them. Creating such inclusive environments is pivotal to realizing the full potential of AR and empowering diverse communities to engage with AR tools effectively.

6.4.3 Personalised yet Collective & Communal

Collaboration and a sense of community play a pivotal role in meaningful work and can serve as a catalyst to overcome initial challenges related to access and practicality of AR tools. Moreover, fostering a collaborative and communal approach should be a fundamental aspect of working in blended reality environments. While Augmented and Virtual Reality technology holds immense potential, there is also a risk of exacerbating the hyper-individualistic and isolating tendencies of the modern era, especially in the age of the internet, where echo chambers and tailored digital environments prevail. To counteract this risk, it is essential to strive for a balance between personalisation and collective engagement.

Physical tools are a potential countermeasure by acting as a shared "handle" within the virtual environment, bridging the gap between users even when their display mediums differ, ranging from mobile devices to head-mounted AR/VR or even no display at all. These physical tools provide a shared means of mediation, fostering collaboration and understanding despite varying perspectives on the virtual world. Furthermore, just as physical tools and objects in our everyday environments are instrumental in organizing labor [218], the same holds true for virtual environments. Physical tools designed specifically for AR can incorporate collaboration as a foundational principle, enabling multiple users to operate a shared tool simultaneously or utilising physical modifications to distribute a tool among users. By considering collaborative design from the outset, these physical tools can enhance the collaborative potential of AR, encouraging joint engagement and enabling a collective approach to problem-solving and task execution.

Moreover, physical tools not only enable collaboration among collocated users but also facilitate work between remote users. A prime example from this work, is how containers in the cubic tools design space can act as 'portals' to pass files, media, or virtual objects between remote users in remote virtual environments. While this thesis primarily focused on user-actuated tools with passive haptics, it is worth considering computer-actuated tools with active haptics, as they hold potential applications for remote collaboration in virtual environments. In such scenarios, a tool operated by a local user could dynamically modify itself based on explicit cues, the contextual requirements of the work, or instructions from a remote user. This adaptability allows for real-time interaction and mutual influence between remote users, facilitating effective collaboration in virtual environments regardless of physical distance.

Similarly, when it comes to tool-making and modification in virtual environments, it is intuitive for individuals to collaborate with others, drawing upon different areas of expertise while retaining control and autonomy over their respective tools. Even in the act of using tools and collaborating within virtual environments, there exists a delicate equilibrium between individual ownership and autonomy, and collective participation towards a shared activity or goal, mirroring real-world dynamics [218, 224, 336]. It is important to acknowledge that collaborative AR and VR research had a lesser emphasis in this work, being less of a focus of the empirical studies described in prior chapters. However, moving forward, this body of literature should occupy a central role in the continued exploration and research of AR tools. By actively incorporating collaborative AR and VR research, we can better understand and develop the frameworks necessary to enable effective collaboration, shared tool-making, and collective problem-solving in virtual environments.

6.4.4 Summary

In summary, the key aspects that distinguish AR tools from AR devices are as follows:

1. Integration of Making, Modifying, and Use: The process of making, modifying, and using AR tools should be viewed as a unified process of adoption. Users

should be empowered to take on roles as designers and researchers within virtual environments, enabling them to shape their own work experiences.

2. Heterogeneous and Ephemeral Nature: AR tools should be diverse and heterogeneous, with the capacity to be both physically and virtually ephemeral. They are transient and adaptable, molded by the specific activities and contextual demands of the work being performed.

3. Collaborative and Shared Design: Unlike devices, AR tools are not solely designed for individual use. They should be inherently collaborative in nature and intended to be shared among users. Collaboration is a fundamental aspect of their design and functionality.

By acknowledging these defining characteristics, we can foster the development of AR tools that empower users, facilitate dynamic and context-driven work, and encourage collaboration and shared experiences within virtual environments. Embracing these principles paves the way for the creation of innovative and transformative Augmented Reality experiences.

Chapter 7

Conclusion

In this thesis, we explored approaches for making, modifying, and using physical tools in AR and demonstrate the broad capabilities of interaction facilitated by physically and virtually-modifiable tools. For *Tool-Making*, we developed a framework for classifying physically-modifiable objects in terms of their input capacity. Through two empirical studies inspired by research-through-design and guessability methodologies, we analyzed existing exemplars of physically-modifiable objects and explored how object properties can be linked to interactions in AR. For *Tool-Modifying*, we created a toolkit of interactive cubic tools that can detect user touch and be tracked in 3D. We evaluated the toolkit through various demonstrations and interactions, showcasing its capabilities. For *Tool-Using*, we conducted three empirical studies. Two of the studies focussed on the design and behavioural phenomena of a precise 3D manipulation technique using cubic tools, while the third study compared this technique to state-of-the-art VR/AR hand-controller approaches. As a result, there are several contributions that advance our understanding of physically and virtually modifiable tools in AR and their impact on user interaction and experience:

1. A conceptual framework for physically-modifiable objects in AR, enabling their characterization and comparison based on their distinct properties and their potential to be leveraged for input.

2. New empirical insights from user-elicited design on the suitability of different physically-modifiable object properties for various AR activities.

3. A toolkit for rapid fabrication and prototyping of physically-modifiable, interactive cubic tools, along with a comprehensive design space showcasing different interaction metaphors and techniques using cubic tools in AR.

4. Insights on a precise 3D manipulation technique using cubic tools, which aids in the design of virtual modification and characterizes their overall performance.

> These findings highlight the complex interplay between handedness and the cubic tools and its effect on users' cognitive and physical labour. Additionally, we gain an initial understanding of how simple form factors, such as cubes, compare with state-of-the-art techniques in AR.

Overall, this thesis proposes that as physical and virtual environments become increasingly blended, the adoption of *physical tools* is a central element to address existing interaction challenges and facilitating future labour. The work has demonstrated how new forms of physical AR tools create new types of interaction and overcome the traditional constraints of Tangible interfaces and physical devices, which lack versatility compared to digital tools, which can be achieved through physical and virtual modification.

In the previous chapter, we discussed the nuances of tool modification for interaction. We also highlighted the central role that the physical environment plays in tool-making and using. We noted that physical characteristics of objects have a collective understanding of applicability, but individual designs are often disputed. The discussion concludes with reflections on the potential directions of AR and Tools. It distinguishes AR Tools from current physical interaction devices, and puts forth a perspective and position on how work in virtual environments and AR Tools should be treated in the future.

We exist in a physical world and are called to adapt to a number of different environments where we engage in practical, procedural, and creative work. If the vision of pervasive and nomadic AR is to be realised, our instruments of virtual labour must also adapt to the environment with us; from surface-supported work, as explored in this thesis, to broader domestic and industrial settings. It is interesting to contemplate what tools will look like in the future as virtual and physical work converge. Designers and AR toolmakers alike must take into account the physical environment when ideating and constructing AR tools. In the physical world, our tools are deeply personal and personalised to our individual skills and abilities, something which should be made readily accessible in AR tools through their virtual and physical-modification.

References

[1] *12-Sided Fidget*. URL: https://www.amazon.co.uk/VCOSTORE-Decompression-Sensory-Relieves-Anxiety/dp/B08B3NDT5L/ref=sr_1_1_sspa?keywords=12+sided+fidget+cube&qid=1676295763&sr=8-1-spons&sp_csd=d2lkZ2V0TmFtZT1zcF9hdGY&psc=1. (accessed 15/01/2023).

[2] Nur Ameerah Binti Abdul Halim and Ajune Wanis Binti Ismail. "Raycasting method using hand gesture for target selection on the occluded object in handheld Augmented Reality". *2021 6th IEEE International Conference on Recent Advances and Innovations in Engineering (ICRAIE)*. Vol. 6. 2021, pp. 1–6. DOI: 10.1109/ICRAIE52900.2021.9704035.

[3] Muhammad Abdullah, Minji Kim, Waseem Hassan, Yoshihiro Kuroda, and Seokhee Jeon. "HapticDrone: An Encountered-Type Kinesthetic Haptic Interface with Controllable Force Feedback: Initial Example for 1D Haptic Feedback". *Adjunct Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. UIST '17 Adjunct. Québec City, QC, Canada: Association for Computing Machinery, 2017, pp. 115–117. ISBN: 9781450354196. DOI: 10.1145/3131785.3131821. URL: https://doi.org/10.1145/3131785.3131821.

[4] Parastoo Abtahi and Sean Follmer. "Visuo-Haptic Illusions for Improving the Perceived Performance of Shape Displays". *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: Association for Computing Machinery, 2018, pp. 1–13. ISBN: 9781450356206. DOI: 10.1145/3173574.3173724. URL: https://doi.org/10.1145/3173574.3173724.

[5] Parastoo Abtahi, Sidney Q. Hough, James A. Landay, and Sean Follmer. "Beyond Being Real: A Sensorimotor Control Perspective on Interactions in Virtual Reality". *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI '22. New Orleans, LA, USA: Association for Computing Machinery, 2022. ISBN: 9781450391573. DOI: 10.1145/3491102.3517706. URL: https://doi.org/10.1145/3491102.3517706.

[6] Syed Shah Alam, Samiha Susmit, Chieh-Yu Lin, Mohammad Masukujjaman, and Yi-Hui Ho. "Factors Affecting Augmented Reality Adoption in the Retail Industry". *Journal of Open Innovation: Technology, Market, and Complexity* 7.2 (2021). ISSN: 2199-8531. DOI: 10.3390/joitmc7020142. URL: https://www.mdpi.com/2199-8531/7/2/142.

[7] Pär Anders Albinsson and Shumin Zhai. "High precision touch screen interaction". *Conference on Human Factors in Computing Systems - Proceedings*. 2003.

[8] Jason Alexander, Anne Roudaut, Jürgen Steimle, Kasper Hornbæk, Miguel Bruns Alonso, Sean Follmer, and Timothy Merritt. "Grand Challenges in Shape-Changing Interface Research". *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: Association for Computing Machinery, 2018, pp. 1–14. ISBN: 9781450356206. DOI: 10.1145/3173574.3173873. URL: https://doi.org/10.1145/3173574.3173873.

[9] Mona Alkhattabi. "Augmented Reality as E-learning Tool in Primary Schools' Education: Barriers to Teachers' Adoption." *International Journal of Emerging Technologies in Learning* 12.2 (2017).

[10] Leonardo Angelini, Denis Lalanne, Elise Van den Hoven, Omar Abou Khaled, and Elena Mugellini. "Move, Hold and Touch: A Framework for Tangible Gesture Interactive Systems". *Machines* 3.3 (2015), pp. 173–207. ISSN: 2075-1702. DOI: 10.3390/machines3030173. URL: https://www.mdpi.com/2075-1702/3/3/173.

[11] Marian Annett. "Hand preference observed in large healthy samples: Classification, norms and interpretations of increased non-right-handedness by the right shift theory". *British journal of Psychology* 95.3 (2004), pp. 339–353.

[12] Alissa N Antle and Sijie Wang. "Comparing motor-cognitive strategies for spatial problem solving with tangible and multi-touch interfaces". *Proceedings of the 7th International Conference on tangible, embedded and embodied interaction*. 2013, pp. 65–72.

[13] Bruno Araujo, Ricardo Jota, Varun Perumal, Jia Xian Yao, Karan Singh, and Daniel Wigdor. "Snake Charmer: Physically Enabling Virtual Objects". *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '16. Eindhoven, Netherlands: Association for Computing Machinery, 2016, pp. 218–226. ISBN: 9781450335829. DOI: 10.1145/2839462.2839484. URL: https://doi.org/10.1145/2839462.2839484.

[14] Ferran Argelaguet and Carlos Andujar. "A survey of 3D object selection techniques for virtual environments". *Computers and Graphics (Pergamon)* (2013). ISSN: 00978493. DOI: 10.1016/j.cag.2012.12.003.

[15] Ahmed Sabbir Arif, Roozbeh Manshaei, Sean Delong, Brien East, Matthew Kyan, and Ali Mazalek. "Sparse Tangibles: Collaborative exploration of gene networks using active tangibles and interactive tabletops". *Proceedings of the TEI'16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*. 2016, pp. 287–295.

[16] Ryan Arisandi, Mai Otsuki, Asako Kimura, Fumihisa Shibata, and Hideyuki Tamura. "Virtual Handcrafting: Building Virtual Wood Models Using ToolDevice". *Proceedings of the IEEE* 102.2 (2014), pp. 185–195. DOI: 10.1109/JPROC.2013.2294243.

[17] Jatin Arora, Varnit Jain, Aryan Saini, Shwetank Shrey, Nirmita Mehra, and Aman Parnami. "VirtualBricks: Exploring a scalable, modular toolkit for enabling physical manipulation in VR". *Conference on Human Factors in Computing Systems - Proceedings*. 2019. ISBN: 9781450359702. DOI: 10.1145/3290605.3300286.

[18] Rahul Arora, Rubaiat Habib Kazi, Tovi Grossman, George Fitzmaurice, and Karan Singh. "SymbiosisSketch: Combining 2D &; 3D Sketching for Designing Detailed 3D Objects in Situ". *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery, 2018, pp. 1–15. ISBN: 9781450356206. URL: https://doi.org/10.1145/3173574.3173759.

[19] Daniel Avrahami, Jacob O. Wobbrock, and Shahram Izadi. "Portico: Tangible Interaction on and around a Tablet". *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*. UIST '11. Santa Barbara, California, USA: Association for Computing Machinery, 2011, pp. 347–356. ISBN: 9781450307161. DOI: 10.1145/2047196.2047241. URL: https://doi.org/10.1145/2047196.2047241.

[20] Takumi Azai, Shuhei Ogawa, Mai Otsuki, Fumihisa Shibata, and Asako Kimura. "Selection and Manipulation Methods for a Menu Widget on the Human Forearm". *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. CHI EA '17. Denver, Colorado, USA: Association for Computing Machinery, 2017, pp. 357–360. ISBN: 9781450346566. DOI: 10.1145/3027063.3052959. URL: https://doi.org/10.1145/3027063.3052959.

[21] Mahdi Azmandian, Mark Hancock, Hrvoje Benko, Eyal Ofek, and Andrew D Wilson. "Haptic retargeting: Dynamic repurposing of passive haptics for enhanced virtual reality experiences". *Proceedings of the 2016 chi conference on human factors in computing systems*. 2016, pp. 1968–1979.

[22] Ronald Azuma and Gary Bishop. "Improving Static and Dynamic Registration in an Optical See-through HMD". *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '94. New York, NY, USA: Association for Computing Machinery, 1994, pp. 197–204. ISBN: 0897916670. DOI: 10.1145/192161.192199. URL: https://doi.org/10.1145/192161.192199.

[23] Ronald T Azuma. "A survey of augmented reality". *Presence: teleoperators & virtual environments* 6.4 (1997), pp. 355–385.

[24] Christopher Baber. *Cognition and tool use: Forms of engagement in human and animal use of tools*. CRC Press, 2003.

[25] Myroslav Bachynskyi, Gregorio Palmas, Antti Oulasvirta, Jürgen Steimle, and Tino Weinkauf. "Performance and Ergonomics of Touch Surfaces: A Comparative Study Using Biomechanical Simulation". *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. Seoul, Republic of Korea: Association for Computing Machinery, 2015, pp. 1817–1826. ISBN: 9781450331456. DOI: 10.1145/2702123.2702607. URL: https://doi.org/10.1145/2702123.2702607.

[26] Leia B Bagesteiro and Robert L Sainburg. "Nondominant arm advantages in load compensation during rapid elbow joint movements". *Journal of neurophysiology* 90.3 (2003), pp. 1503–1513.

[27] Till Ballendat, Nicolai Marquardt, and Saul Greenberg. "Proxemic interaction: designing for a proximity and orientation-aware environment". *ACM International Conference on Interactive Tabletops and Surfaces*. 2010, pp. 121–130.

[28] Marc Baloup, Thomas Pietrzak, and Géry Casiez. "Raycursor: A 3D Pointing Facilitation Technique based on Raycasting". *Conference on Human Factors in Computing Systems - Proceedings*. 2019. ISBN: 9781450359702. DOI: 10.1145/3290605.3300331.

[29] Reigo Ban, Keigo Matsumoto, Takuji Narumi, and Hideaki Kuzuoka. "Wormholes in VR: Teleporting Hands for Flexible Passive Haptics". *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 2022, pp. 748–757. DOI: 10.1109/ISMAR55827.2022.00093.

[30] Patrick Baudisch. "Virtual Reality in Your Living Room: Technical Perspective". *Commun. ACM* 58.6 (May 2015), p. 92. ISSN: 0001-0782. DOI: 10.1145/2754393. URL: https://doi.org/10.1145/2754393.

[31] Michel Beaudouin-Lafon. "Instrumental Interaction: An Interaction Model for Designing Post-WIMP User Interfaces". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '00. The Hague, The Netherlands: Association for Computing Machinery, 2000, pp. 446–453. ISBN: 1581132166. DOI: 10.1145/332040.332473. URL: https://doi.org/10.1145/332040.332473.

[32] Michel Beaudouin-Lafon and Wendy E Mackay. "Prototyping tools and techniques". *Human-Computer Interaction*. CRC Press, 2009, pp. 137–160.

[33] Hrvoje Benko, Christian Holz, Mike Sinclair, and Eyal Ofek. "NormalTouch and TextureTouch: High-Fidelity 3D Haptic Shape Rendering on Handheld Virtual Reality Controllers". *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. UIST '16. Tokyo, Japan: Association for Computing Machinery, 2016, pp. 717–728. ISBN: 9781450341899. DOI: 10.1145/2984511.2984526. URL: https://doi.org/10.1145/2984511.2984526.

[34] Hrvoje Benko, Ricardo Jota, and Andrew Wilson. "MirageTable: Freehand Interaction on a Projected Augmented Reality Tabletop". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: Association for Computing Machinery, 2012, pp. 199–208. ISBN: 9781450310154. DOI: 10.1145/2207676.2207704. URL: https://doi.org/10.1145/2207676.2207704.

[35] François Bérard, Jessica Ip, Mitchel Benovoy, Dalia El-Shimy, Jeffrey R. Blum, and Jeremy R. Cooperstock. "Did "Minority Report" Get it Wrong? Superiority of the Mouse over 3D Input in a 3D Placement Task". *Conference Proceedings of INTERACT '09, the Twelfth IFIP TC13 International Conference on Human-Computer Interaction, Springer LNCS (Lecture Notes in Computer Science)*. Berlin, Heidelberg, 2009, pp. 400–414. URL: https://hal.inria.fr/hal-00953677.

[36] Oriel Bergig, Eyal Soreq, Nate Hagbi, Kirill Pevzner, Nati Levi, Shoham Blau, Yulia Smelansky, and Jihad El-Sana. "Out of the cube: augmented Rubik's cube". *International Journal of Computer Games Technology* 2011 (2011).

[37] Joanna Bergström and Kasper Hornbæk. "Human–Computer Interaction on the Skin". *ACM Comput. Surv.* 52.4 (Aug. 2019). ISSN: 0360-0300. DOI: 10.1145/3332166. URL: https://doi.org/10.1145/3332166.

[38] Lonni Besançon, Paul Issartel, Mehdi Ammi, and Tobias Isenberg. "Mouse, tactile, and tangible input for 3D manipulation". *Conference on Human Factors in Computing Systems - Proceedings*. 2017. ISBN: 9781450346559. DOI: 10.1145/3025453.3025863.

[39] Eric A. Bier. "Snap-Dragging in Three Dimensions". *Proceedings of the 1990 Symposium on Interactive 3D Graphics*. I3D '90. Snowbird, Utah, USA: Association for Computing Machinery, 1990, pp. 193–204. ISBN: 0897913515. DOI: 10.1145/91385.91446. URL: https://doi.org/10.1145/91385.91446.

[40] Eric A. Bier. "Snap-Dragging in Three Dimensions". *Proceedings of the 1990 Symposium on Interactive 3D Graphics*. I3D '90. Snowbird, Utah, USA: Association for Computing Machinery, 1990, pp. 193–204. ISBN: 0897913515. DOI: 10.1145/91385.91446. URL: https://doi.org/10.1145/91385.91446.

[41] Mark Billinghurst. "Shared space: Collaborative augmented reality". *ACM SIGGRAPH 99 Conference abstracts and applications*. 1999, p. 178.

[42] Mark Billinghurst. "Augmented reality in education". *New horizons for learning* 12.5 (2002), pp. 1–5.

[43] Mark Billinghurst, Adrian Clark, and Gun Lee. "A Survey of Augmented Reality". *Foundations and Trends® in Human–Computer Interaction* 8.2-3 (2015), pp. 73–272. ISSN: 1551-3955. DOI: 10.1561/1100000049. URL: http://dx.doi.org/10.1561/1100000049.

[44] Mark Billinghurst, Raphael Grasset, and Julian Looser. "Designing augmented reality interfaces". *ACM Siggraph Computer Graphics* 39.1 (2005), pp. 17–22.

[45] Mark Billinghurst, Hirokazu Kato, and Ivan Poupyrev. "Collaboration with tangible augmented reality interfaces". *HCI international*. Vol. 1. 2001, pp. 5–10.

[46] Mark Billinghurst, Hirokazu Kato, Ivan Poupyrev, et al. "Tangible augmented reality". *Acm siggraph asia* 7.2 (2008), pp. 1–10.

[47] Mark Billinghurst, Suzanne Weghorst, and T Furness. "Shared space: An augmented reality approach for computer supported collaborative work". *Virtual Reality* 3 (1998), pp. 25–36.

[48] Mark Billinghurst, Suzanne Weghorst, and Tom Furness. "Wearable computers for three dimensional CSCW". *Digest of Papers. First International Symposium on Wearable Computers*. IEEE. 1997, pp. 39–46.

[49] Oliver Bimber and Ramesh Raskar. *Spatial augmented reality: merging real and virtual worlds*. CRC press, 2005.

[50] Renaud Blanch, Yves Guiard, and Michel Beaudouin-Lafon. "Semantic pointing: Improving target acquisition with control-display ratio adaptation". *Conference on Human Factors in Computing Systems - Proceedings*. 2004.

[51] *Blender*. URL: https://www.blender.org/. (accessed 20/01/2023).

[52] YAGV Boas. “Overview of virtual reality technologies”. *Interactive Multimedia Conference*. Vol. 2013. 2013.

[53] Susanne Bødker. *Through the Interface: A Human Activity Approach to User Interface Design*. USA: L. Erlbaum Associates Inc., 1990. ISBN: 0805805710.

[54] Doug A Bowman and Larry F Hodges. “An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments”. *Proceedings of the 1997 symposium on Interactive 3D graphics*. 1997, 35–ff.

[55] Doug A. Bowman, Sabine Coquillart, Bernd Froehlich, Michitaka Hirose, Yoshifumi Kitamura, Kiyoshi Kiyokawa, and Wolfgang Stuerzlinger. “3D User Interfaces: New Directions and Perspectives”. *IEEE Computer Graphics and Applications* 28.6 (2008), pp. 20–36. DOI: 10.1109/MCG.2008.109.

[56] Sean Braley, Calvin Rubens, Timothy R. Merritt, and Roel Vertegaal. “GridDrones: A Self-Levitating Physical Voxel Lattice for 3D Surface Deformations”. *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI EA ’18. Montreal QC, Canada: Association for Computing Machinery, 2018, pp. 1–4. ISBN: 9781450356213. DOI: 10.1145/3170427.3186477. URL: https://doi.org/10.1145/3170427.3186477.

[57] Virginia Braun and Victoria Clarke. “Using thematic analysis in psychology”. *Qualitative research in psychology* 3.2 (2006), pp. 77–101.

[58] Brenda Brouwer, Martin V Sale, and Michael A Nordstrom. “Asymmetry of motor cortex excitability during a simple motor task: relationships with handedness and manual performance”. *Experimental Brain Research* 138.4 (2001), pp. 467–476.

[59] Frederik Brudy, Christian Holz, Roman Rädle, Chi-Jui Wu, Steven Houben, Clemens Nylandsted Klokmose, and Nicolai Marquardt. “Cross-Device Taxonomy: Survey, Opportunities and Challenges of Interactions Spanning Across Multiple Devices”. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI ’19. Glasgow, Scotland Uk: Association for Computing Machinery, 2019, pp. 1–28. ISBN: 9781450359702. DOI: 10.1145/3290605.3300792. URL: https://doi.org/10.1145/3290605.3300792.

[60] Volkert Buchmann, Stephen Violich, Mark Billinghurst, and Andy Cockburn. “FingARtips: Gesture Based Direct Manipulation in Augmented Reality”. *Proceedings of the 2nd International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia*. GRAPHITE ’04. Singapore: Association for Computing Machinery, 2004, pp. 212–221. ISBN: 1581138830. DOI: 10.1145/988834.988871. URL: https://doi.org/10.1145/988834.988871.

[61] Jesse Burstyn, Nicholas Fellion, Paul Strohmeier, and Roel Vertegaal. “Printput: Resistive and capacitive input widgets for interactive 3D prints”. *Human-Computer Interaction–INTERACT 2015: 15th IFIP TC 13 International Conference, Bamberg, Germany, September 14-18, 2015, Proceedings, Part I 15*. Springer. 2015, pp. 332–339.

[62] Jeff Butterworth, Andrew Davidson, Stephen Hench, and Marc T Olano. “3DM: A three dimensional modeler using a head-mounted display”. *Proceedings of the 1992 symposium on Interactive 3D graphics*. 1992, pp. 135–138.

[63] W. Buxton and B. Myers. “A Study in Two-Handed Input”. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’86. Boston, Massachusetts, USA: Association for Computing Machinery, 1986, pp. 321–326. ISBN: 0897911806. DOI: 10.1145/22627.22390. URL: https://doi.org/10.1145/22627.22390.

[64] Stuart K Card, Thomas P Moran, and Allen Newell. *The psychology of human-computer interaction*. Crc Press, 2018.

[65] Julie Carmigniani, Borko Furht, Marco Anisetti, Paolo Ceravolo, Ernesto Damiani, and Misa Ivkovic. “Augmented reality technologies, systems and applications”. *Multimedia tools and applications* 51.1 (2011), pp. 341–377.

[66] John M Carroll, Wendy A Kellogg, and Mary Beth Rosson. “The task-artifact cycle”. *Designing interaction: Psychology at the human-computer interface* (1991), pp. 74–102.

[67] Tom Carter, Sue Ann Seah, Benjamin Long, Bruce Drinkwater, and Sriram Subramanian. “UltraHaptics: Multi-Point Mid-Air Haptic Feedback for Touch Surfaces”. *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*. UIST ’13. St. Andrews, Scotland, United Kingdom: Association for Computing Machinery, 2013, pp. 505–514. ISBN: 9781450322683. DOI: 10.1145/2501988.2502018. URL: https://doi.org/10.1145/2501988.2502018.

[68] Géry Casiez, Daniel Vogel, Ravin Balakrishnan, and Andy Cockburn. “The impact of control-display gain on user performance in pointing tasks”. *Human-Computer Interaction* (2008). ISSN: 07370024. DOI: 10.1080/07370020802278163.

[69] Simona Casini, Matteo Morvidoni, Matteo Bianchi, Manuel Catalano, Giorgio Grioli, and Antonio Bicchi. “Design and realization of the CUFF - clenching upper-limb force feedback wearable device for distributed mechano-tactile stimulation of normal and tangential skin forces”. *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2015, pp. 1186–1193. DOI: 10.1109/IROS.2015.7353520.

[70] Edwin Chan, Teddy Seyed, Wolfgang Stuerzlinger, Xing Dong Yang, and Frank Maurer. "User elicitation on single-hand microgestures". *Conference on Human Factors in Computing Systems - Proceedings*. 2016. ISBN: 9781450333627. DOI: 10.1145/2858036.2858589.

[71] Liwei Chan, Stefanie Müller, Anne Roudaut, and Patrick Baudisch. "CapStones and ZebraWidgets: sensing stacks of building blocks, dials and sliders on capacitive touch screens". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2012, pp. 2189–2192.

[72] Emmanuelle Chapoulie, Theophanis Tsandilas, Lora Oehlberg, Wendy Mackay, and George Drettakis. "Finger-based manipulation in immersive spaces and the real world". *2015 IEEE Symposium on 3D User Interfaces (3DUI)*. 2015, pp. 109–116. DOI: 10.1109/3DUI.2015.7131734.

[73] Steven K. Charles and Neville Hogan. "Dynamics of wrist rotations". *Journal of Biomechanics* 44.4 (2011), pp. 614–621. ISSN: 0021-9290. DOI: https://doi.org/10.1016/j.jbiomech.2010.11.016. URL: https://www.sciencedirect.com/science/article/pii/S0021929010006408.

[74] Dimitris Chatzopoulos, Carlos Bermejo, Zhanpeng Huang, and Pan Hui. "Mobile Augmented Reality Survey: From Where We Are to Where We Go". *IEEE Access* 5 (2017), pp. 6917–6950. DOI: 10.1109/ACCESS.2017.2698164.

[75] Di Laura Chen, Ravin Balakrishnan, and Tovi Grossman. "Disambiguation Techniques for Freehand Object Manipulations in Virtual Reality". *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 2020, pp. 285–292. DOI: 10.1109/VR46266.2020.00048.

[76] Michael Chen, S. Joy Mountford, and Abigail Sellen. "A Study in Interactive 3-D Rotation Using 2-D Control Devices". *SIGGRAPH Comput. Graph.* 22.4 (June 1988), pp. 121–129. ISSN: 0097-8930. DOI: 10.1145/378456.378497. URL: https://doi.org/10.1145/378456.378497.

[77] Peng Chen, Xiaolin Liu, Wei Cheng, and Ronghuai Huang. "A review of using Augmented Reality in Education from 2011 to 2016". *Innovations in smart learning* (2017), pp. 13–18.

[78] Zhaorui Chen, Jinzhu Li, Yifan Hua, Rui Shen, and Anup Basu. "Multimodal interaction in augmented reality". *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. 2017, pp. 206–209. DOI: 10.1109/SMC.2017.8122603.

[79] Zhen Chen, Xiaochi Ma, Zeya Peng, Ying Zhou, Mengge Yao, Zheng Ma, Ci Wang, Zaifeng Gao, and Mowei Shen. "User-Defined Gestures for Gestural Interaction: Extending from Hands to Other Body Parts". *International Journal of Human-Computer Interaction* (2018). ISSN: 15327590. DOI: 10.1080/10447318.2017.1342943.

[80] Lung-Pan Cheng, Li Chang, Sebastian Marwecki, and Patrick Baudisch. "ITurk: Turning Passive Haptics into Active Haptics by Making Users Reconfigure Props in Virtual Reality". *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: Association for Computing Machinery, 2018, pp. 1–10. ISBN: 9781450356206. DOI: 10.1145/3173574.3173663. URL: https://doi.org/10.1145/3173574.3173663.

[81] Keshav Chintamani, Alex Cao, R. Darin Ellis, and Abhilash K. Pandya. "Improved Telemanipulator Navigation During Display-Control Misalignments Using Augmented Reality Cues". *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 40.1 (2010), pp. 29–39. DOI: 10.1109/TSMCA.2009.2030166.

[82] Inrak Choi, Heather Culbertson, Mark R. Miller, Alex Olwal, and Sean Follmer. "Grabity: A Wearable Haptic Interface for Simulating Weight and Grasping in Virtual Reality". *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. UIST '17. Québec City, QC, Canada: Association for Computing Machinery, 2017, pp. 119–130. ISBN: 9781450349819. DOI: 10.1145/3126594.3126599. URL: https://doi.org/10.1145/3126594.3126599.

[83] Inrak Choi, Elliot W. Hawkes, David L. Christensen, Christopher J. Ploch, and Sean Follmer. "Wolverine: A wearable haptic interface for grasping in virtual reality". *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2016, pp. 986–993. DOI: 10.1109/IROS.2016.7759169.

[84] Inrak Choi, Eyal Ofek, Hrvoje Benko, Mike Sinclair, and Christian Holz. "CLAW: A Multifunctional Handheld Haptic Controller for Grasping, Touching, and Triggering in Virtual Reality". *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: Association for Computing Machinery, 2018, pp. 1–13. ISBN: 9781450356206. DOI: 10.1145/3173574.3174228. URL: https://doi.org/10.1145/3173574.3174228.

[85] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. "Empirical evaluation of gated recurrent neural networks on sequence modeling". *arXiv preprint arXiv:1412.3555* (2014).

[86] Keywon Chung, Michael Shilman, Chris Merrill, and Hiroshi Ishii. "OnObject: gestural play with tagged everyday objects". *Adjunct proceedings of the 23nd annual ACM symposium on user interface software and technology.* 2010, pp. 379–380.

[87] Marcelo Coelho, Hiroshi Ishii, and Pattie Maes. "Surflex: a programmable surface for the design of tangible interfaces". *CHI'08 extended abstracts on Human factors in computing systems.* 2008, pp. 3429–3434.

[88] P. R. Cohen, M. Dalrymple, D. B. Moran, F. C. Pereira, and J. W. Sullivan. "Synergistic Use of Direct Manipulation and Natural Language". *SIGCHI Bull.* 20.SI (Mar. 1989), pp. 227–233. ISSN: 0736-6906. DOI: 10.1145/67450.67494. URL: https://doi.org/10.1145/67450.67494.

[89] Philip R. Cohen, Michael Johnston, David McGee, Sharon Oviatt, Jay Pittman, Ira Smith, Liang Chen, and Josh Clow. "QuickSet: Multimodal Interaction for Distributed Applications". *Proceedings of the Fifth ACM International Conference on Multimedia.* MULTIMEDIA '97. Seattle, Washington, USA: Association for Computing Machinery, 1997, pp. 31–40. ISBN: 0897919912. DOI: 10.1145/266180.266328. URL: https://doi.org/10.1145/266180.266328.

[90] Emmanuel Courtoux, Caroline Appert, and Olivier Chapuis. "WallTokens: Surface Tangibles for Vertical Displays". *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems.* 2021, pp. 1–13.

[91] Céline Coutrix and Laurence Nigay. "Mixed Reality: A Model of Mixed Interaction". *Proceedings of the Working Conference on Advanced Visual Interfaces.* AVI '06. Venezia, Italy: Association for Computing Machinery, 2006, pp. 43–50. ISBN: 1595933530. DOI: 10.1145/1133265.1133274. URL: https://doi.org/10.1145/1133265.1133274.

[92] Carolina Cruz-Neira, Daniel J. Sandin, Thomas A. DeFanti, Robert V. Kenyon, and John C. Hart. "The CAVE: Audio Visual Experience Automatic Virtual Environment". *Commun. ACM* 35.6 (June 1992), pp. 64–72. ISSN: 0001-0782. DOI: 10.1145/129888.129892. URL: https://doi.org/10.1145/129888.129892.

[93] *Cube Spinner.* URL: https://www.multi-sensoryworld.co.uk/products/cube-spinner. (accessed 15/01/2023).

[94] Heather Culbertson, Samuel B Schorr, and Allison M Okamura. "Haptics: The present and future of artificial touch sensation". *Annual Review of Control, Robotics, and Autonomous Systems* 1.1 (2018), pp. 385–409.

[95] *CyberGlove CyberGrasp.* URL: http://www.cyberglovesystems.com/cybergrasp. (accessed 24/01/2023).

[96] Maxime Daniel, Guillaume Rivière, and Nadine Couture. "CairnFORM: A Shape-Changing Ring Chart Notifying Renewable Energy Availability in Peripheral Locations". *Proceedings of the Thirteenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '19. Tempe, Arizona, USA: Association for Computing Machinery, 2019, pp. 275–286. ISBN: 9781450361965. DOI: 10.1145/3294109.3295634. URL: https://doi.org/10.1145/3294109.3295634.

[97] Oscar Danielsson, Magnus Holm, and Anna Syberfeldt. "Augmented reality smart glasses in industrial assembly: Current status and future challenges". *Journal of Industrial Information Integration* 20 (2020), p. 100175. ISSN: 2452-414X. DOI: https://doi.org/10.1016/j.jii.2020.100175. URL: https://www.sciencedirect.com/science/article/pii/S2452414X20300509.

[98] Bruno R. De Araújo, Géry Casiez, Joaquim A. Jorge, and Martin Hachet. "Mockup Builder: 3D modeling on and above the surface". *Computers and Graphics (Pergamon)*. 2013. DOI: 10.1016/j.cag.2012.12.005.

[99] Michael F. Deering. "HoloSketch: A Virtual Reality Sketching/Animation Tool". *ACM Trans. Comput.-Hum. Interact.* 2.3 (Sept. 1995), pp. 220–238. ISSN: 1073-0516. DOI: 10.1145/210079.210087. URL: https://doi.org/10.1145/210079.210087.

[100] Juan Manuel Davila Delgado, Lukumon Oyedele, Thomas Beach, and Peter Demian. "Augmented and Virtual Reality in Construction: Drivers and Limitations for Industry Adoption". *Journal of Construction Engineering and Management* 146.7 (2020), p. 04020079. DOI: 10.1061/(ASCE)CO.1943-7862.0001844. eprint: https://ascelibrary.org/doi/pdf/10.1061/\%28ASCE\%29CO.1943-7862.0001844. URL: https://ascelibrary.org/doi/abs/10.1061/%5C%28ASCE%5C%29CO.1943-7862.0001844.

[101] *DextaRobotics Dexmo*. URL: https://www.dextarobotics.com/en-us. (accessed 24/01/2023).

[102] Alan Dix, Janet Finlay, Gregory D Abowd, and Russell Beale. *Human-computer interaction*. Pearson Education, 2003.

[103] Paul Dourish. *Where the Action Is: The Foundations of Embodied Interaction*. MIT Press, 2004.

[104] Tobias Drey, Michael Montag, Andrea Vogt, Nico Rixen, Tina Seufert, Steffi Zander, Michael Rietzler, and Enrico Rukzio. "Investigating the Effects of Individual Spatial Abilities on Virtual Reality Object Manipulation". *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. CHI '23. Hamburg, Germany: Association for Computing Machinery,

2023. ISBN: 9781450394215. DOI: 10.1145/3544548.3581004. URL: https://doi.org/10.1145/3544548.3581004.

[105] Adam Drogemuller, James Walsh, Ross T. Smith, Matt Adcock, and Bruce H Thomas. "Turning Everyday Objects into Passive Tangible Controllers". *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '21. Salzburg, Austria: Association for Computing Machinery, 2021. ISBN: 9781450382137. DOI: 10.1145/3430524.3442460. URL: https://doi.org/10.1145/3430524.3442460.

[106] Ruofei Du, Alex Olwal, Mathieu Le Goc, Shengzhi Wu, Danhang Tang, Yinda Zhang, Jun Zhang, David Joseph Tan, Federico Tombari, and David Kim. "Opportunistic Interfaces for Augmented Reality: Transforming Everyday Objects into Tangible 6DoF Interfaces Using Ad Hoc UI". *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI EA '22. New Orleans, LA, USA: Association for Computing Machinery, 2022. ISBN: 9781450391566. DOI: 10.1145/3491101.3519911. URL: https://doi.org/10.1145/3491101.3519911.

[107] Brien East, Sean DeLong, Roozbeh Manshaei, Ahmed Arif, and Ali Mazalek. "Actibles: open source active tangibles". *Proceedings of the 2016 ACM International Conference on Interactive Surfaces and Spaces*. 2016, pp. 469–472.

[108] S.R. Ellis. "What are virtual environments?" *IEEE Computer Graphics and Applications* 14.1 (1994), pp. 17–22. DOI: 10.1109/38.250914.

[109] Douglas C Engelbart and William K English. "A research center for augmenting human intellect". *Proceedings of the December 9-11, 1968, fall joint computer conference, part I*. 1968, pp. 395–410.

[110] Severin Engert, Konstantin Klamka, Andreas Peetz, and Raimund Dachselt. "STRAIDE: A Research Platform for Shape-Changing Spatial Displays Based on Actuated Strings". *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI '22. New Orleans, LA, USA: Association for Computing Machinery, 2022. ISBN: 9781450391573. DOI: 10.1145/3491102.3517462. URL: https://doi.org/10.1145/3491102.3517462.

[111] David Englmeier, Julia Dörner, Andreas Butz, and Tobias Höllerer. "A tangible spherical proxy for object manipulation in augmented reality". *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE. 2020, pp. 221–229.

[112] Marc O Ernst and Martin S Banks. "Humans integrate visual and haptic information in a statistically optimal fashion". *Nature* 415.6870 (2002), pp. 429–433.

[113] Marc O. Ernst and Heinrich H. BÃ¼lthoff. “Merging the senses into a robust percept”. *Trends in Cognitive Sciences* 8.4 (2004), pp. 162–169. ISSN: 1364-6613. DOI: https://doi.org/10.1016/j.tics.2004.02.002. URL: https://www.sciencedirect.com/science/article/pii/S1364661304000385.

[114] Daniel Fallman. “Design-Oriented Human-Computer Interaction”. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’03. Ft. Lauderdale, Florida, USA: Association for Computing Machinery, 2003, pp. 225–232. ISBN: 1581136307. DOI: 10.1145/642611.642652. URL: https://doi.org/10.1145/642611.642652.

[115] Simone Fani, Simone Ciotti, Edoardo Battaglia, Alessandro Moscatelli, and Matteo Bianchi. “W-FYD: A Wearable Fabric-Based Display for Haptic Multi-Cue Delivery and Tactile Augmented Reality”. *IEEE Transactions on Haptics* 11.2 (2018), pp. 304–316. DOI: 10.1109/TOH.2017.2708717.

[116] Daniel Feeney, Igor Jelaska, Mehmet Uygur, and Slobodan Jaric. “Effects of Unilateral Muscle Fatigue on Performance and Force Coordination in Bimanual Manipulation Tasks”. *Motor Control* 21.1 (2017), pp. 26–41. DOI: 10.1123/mc.2015-0037. URL: https://journals.humankinetics.com/view/journals/mcj/21/1/article-p26.xml.

[117] Martin Feick, Scott Bateman, Anthony Tang, André Miede, and Nicolai Marquardt. “TanGi: Tangible Proxies for Embodied Object Exploration and Manipulation in Virtual Reality”. *arXiv preprint arXiv:2001.03021* (2020).

[118] Alex Olwal Steven Feiner. “The flexible pointer: An interaction technique for selection in augmented and virtual reality”. *Proc. UIST*. Vol. 3. 2003, pp. 81–82.

[119] Ylva Fernaeus, Jakob Tholander, and Martin Jonsson. “Towards a new set of ideals: consequences of the practice turn in tangible interaction”. *Proceedings of the 2nd international conference on Tangible and embedded interaction*. 2008, pp. 223–230.

[120] Tiare Feuchtner and Jörg Müller. “Extending the Body for Interaction with Reality”. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI ’17. Denver, Colorado, USA: Association for Computing Machinery, 2017, pp. 5145–5157. ISBN: 9781450346559. DOI: 10.1145/3025453.3025689. URL: https://doi.org/10.1145/3025453.3025689.

[121] M. Fiala. “ARTag, a fiducial marker system using digital techniques”. *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*. Vol. 2. 2005, 590–596 vol. 2. DOI: 10.1109/CVPR.2005.74.

[122] *Fiddle Stick*. URL: https://ellusionist.com/products/fiddle-stick. (accessed 15/01/2023).

[123] *Fidget Cube*. URL: https://www.antsylabs.com/collections/fidget/products/fidget-cube-the-original-fidget-toy. (accessed 15/01/2023).

[124] *Finger Trainer*. URL: https://www.amazon.co.uk/Yakamoz-Portable-Exerciser-Trainer-Strengthener/dp/B01G6TRMHW/ref=pd_lpo_1?pd_rd_w=OpUW3&content-id=amzn1.sym.2d229339-2f42-4596-a90d-b81a4f52d6d3&pf_rd_p=2d229339-2f42-4596-a90d-b81a4f52d6d3&pf_rd_r=TBAD0QEGPN3NRX36ECQD&pd_rd_wg=QRwSr&pd_rd_r=897eace6-77b4-4c09-b5ff-5456031e7d3c&pd_rd_i=B01G6TRMHW&psc=1. (accessed 15/01/2023).

[125] Kenneth P Fishkin. "A taxonomy for and analysis of tangible interfaces". *Personal and Ubiquitous computing* 8.5 (2004), pp. 347–358.

[126] Kenneth P Fishkin, Thomas P Moran, and Beverly L Harrison. "Embodied user interfaces: Towards invisible user interfaces". *IFIP International Conference on Engineering for Human-Computer Interaction*. Springer. 1998, pp. 1–18.

[127] Paul M Fitts. "The information capacity of the human motor system in controlling the amplitude of movement." *Journal of experimental psychology* 47.6 (1954), p. 381.

[128] George W Fitzmaurice, Hiroshi Ishii, and William AS Buxton. "Bricks: laying the foundations for graspable user interfaces". *Proceedings of the SIGCHI conference on Human factors in computing systems*. 1995, pp. 442–449.

[129] Morten Fjeld, Jonas Fredriksson, Martin Ejdestig, Florin Duca, Kristina Bötschi, Benedikt Voegtli, and Patrick Juchli. "Tangible user interface for chemistry education: comparative evaluation and re-design". *Proceedings of the SIGCHI conference on Human factors in computing systems*. 2007, pp. 805–808.

[130] Sean Follmer, Daniel Leithinger, Alex Olwal, Akimitsu Hogge, and Hiroshi Ishii. "inFORM: dynamic physical affordances and constraints through shape and object actuation." *Uist*. Vol. 13. 10. Citeseer. 2013, pp. 2501–988.

[131] Scott Frees, G. Drew Kessler, and Edwin Kay. "PRISM Interaction for Enhancing Control in Immersive Virtual Environments". *ACM Trans. Comput.-Hum. Interact.* 14.1 (May 2007), 2–es. ISSN: 1073-0516. DOI: 10.1145/1229855.1229857. URL: https://doi.org/10.1145/1229855.1229857.

[132] Maite Frutos-Pascual, Chris Creed, and Ian Williams. "Head mounted display interaction evaluation: manipulating virtual objects in augmented reality". *IFIP Conference on Human-Computer Interaction*. Springer. 2019, pp. 287–308.

[133] David T. Gallant, Andrew G. Seniuk, and Roel Vertegaal. "Towards More Paper-like Input: Flexible Input Devices for Foldable Interaction Styles". *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology*. UIST '08. Monterey, CA, USA: Association for Computing Machinery, 2008, pp. 283–286. ISBN: 9781595939753. DOI: 10.1145/1449715.1449762. URL: https://doi.org/10.1145/1449715.1449762.

[134] Yahya Ghazwani and Shamus Smith. "Interaction in Augmented Reality: Challenges to Enhance User Experience". *Proceedings of the 2020 4th International Conference on Virtual and Augmented Reality Simulations*. ICVARS 2020. Sydney, NSW, Australia: Association for Computing Machinery, 2020, pp. 39–44. ISBN: 9781450376945. DOI: 10.1145/3385378.3385384. URL: https://doi.org/10.1145/3385378.3385384.

[135] Yahya Ghazwani and Shamus Smith. "Interaction in Augmented Reality: Challenges to Enhance User Experience". *Proceedings of the 2020 4th International Conference on Virtual and Augmented Reality Simulations*. ICVARS 2020. Sydney, NSW, Australia: Association for Computing Machinery, 2020, pp. 39–44. ISBN: 9781450376945. DOI: 10.1145/3385378.3385384. URL: https://doi.org/10.1145/3385378.3385384.

[136] James J Gibson. "The theory of affordances". *Hilldale, USA* 1.2 (1977), pp. 67–82.

[137] Antonio Gomes, Calvin Rubens, Sean Braley, and Roel Vertegaal. "BitDrones: Towards Using 3D Nanocopter Displays as Interactive Self-Levitating Programmable Matter". *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. San Jose, California, USA: Association for Computing Machinery, 2016, pp. 770–780. ISBN: 9781450333627. DOI: 10.1145/2858036.2858519. URL: https://doi.org/10.1145/2858036.2858519.

[138] Jens Emil Grønbæk, Mille Skovhus Knudsen, Kenton O'Hara, Peter Gall Krogh, Jo Vermeulen, and Marianne Graves Petersen. "Proxemics beyond proximity: Designing for flexible social interaction through cross-device interaction". *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 2020, pp. 1–14.

[139] Tobias Grosse-Puppendahl, Yannick Berghoefer, Andreas Braun, Raphael Wimmer, and Arjan Kuijper. “OpenCapSense: A rapid prototyping toolkit for pervasive interaction using capacitive sensing”. *2013 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. 2013, pp. 152–159. DOI: 10.1109/PerCom.2013.6526726.

[140] Tobias Grosse-Puppendahl, Christian Holz, Gabe Cohn, Raphael Wimmer, Oskar Bechtold, Steve Hodges, Matthew S. Reynolds, and Joshua R. Smith. “Finding Common Ground: A Survey of Capacitive Sensing in Human-Computer Interaction”. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI ’17. Denver, Colorado, USA: Association for Computing Machinery, 2017, pp. 3293–3315. ISBN: 9781450346559. DOI: 10.1145/3025453.3025808. URL: https://doi.org/10.1145/3025453.3025808.

[141] Tovi Grossman and Ravin Balakrishnan. “The Bubble Cursor: Enhancing Target Acquisition by Dynamic Resizing of the Cursor’s Activation Area”. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI ’05. Portland, Oregon, USA: Association for Computing Machinery, 2005, pp. 281–290. ISBN: 1581139985. DOI: 10.1145/1054972.1055012. URL: https://doi.org/10.1145/1054972.1055012.

[142] Jens Grubert, Tobias Langlotz, Stefanie Zollmann, and Holger Regenbrecht. “Towards Pervasive Augmented Reality: Context-Awareness in Augmented Reality”. *IEEE Transactions on Visualization and Computer Graphics* 23.6 (2017), pp. 1706–1724. DOI: 10.1109/TVCG.2016.2543720.

[143] Jens Grubert, Eyal Ofek, Michel Pahud, and Per Ola Kristensson. “The Office of the Future: Virtual, Portable, and Global”. *IEEE Computer Graphics and Applications* 38.6 (2018), pp. 125–133. DOI: 10.1109/MCG.2018.2875609.

[144] Yves Guiard. “Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model”. *Journal of motor behavior* 19.4 (1987), pp. 486–517.

[145] Aakar Gupta, Cheng Ji, Hui-Shyong Yeo, Aaron Quigley, and Daniel Vogel. “RotoSwype: Word-Gesture Typing Using a Ring”. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI ’19. Glasgow, Scotland Uk: Association for Computing Machinery, 2019, pp. 1–12. ISBN: 9781450359702. DOI: 10.1145/3290605.3300244. URL: https://doi.org/10.1145/3290605.3300244.

[146] Aakar Gupta, Bo Rui Lin, Siyi Ji, Arjav Patel, and Daniel Vogel. "Replicate and Reuse: Tangible Interaction Design for Digitally-Augmented Physical Media Objects". *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 2020, pp. 1–12.

[147] *Gyro Toy*. URL: https://www.amazon.co.uk/PUTOWUT-Illusion-Spherical-Gyroscope-Decompression/dp/B09CMTYDBG/ref=asc_df_B09CMTYDBG/?tag=googshopuk-21&linkCode=df0&hvadid=570437836462&hvpos=&hvnetw=g&hvrand=6026081715041730921&hvpone=&hvptwo=&hvqmt=&hvdev=c&hvdvcmdl=&hvlocint=&hvlocphy=1006854&hvtargid=pla-1629666026527&psc=1. (accessed 15/01/2023).

[148] Changyo Han, Ryo Takahashi, Yuchi Yahagi, and Takeshi Naemura. "PneuModule: Using Inflatable Pin Arrays for Reconfigurable Physical Controls on Pressure-Sensitive Touch Surfaces". *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 2020, pp. 1–14.

[149] Teng Han, Fraser Anderson, Pourang Irani, and Tovi Grossman. "HydroRing: Supporting Mixed Reality Haptics Using Liquid Flow". *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST '18. Berlin, Germany: Association for Computing Machinery, 2018, pp. 913–925. ISBN: 9781450359481. DOI: 10.1145/3242587.3242667. URL: https://doi.org/10.1145/3242587.3242667.

[150] Jeffrey T. Hansberger, Chao Peng, Shannon L. Mathis, Vaidyanath Areyur Shanthakumar, Sarah C. Meacham, Lizhou Cao, and Victoria R. Blakely. "Dispelling the Gorilla Arm Syndrome: The Viability of Prolonged Gesture Interactions". *Virtual, Augmented and Mixed Reality*. Ed. by Stephanie Lackey and Jessie Chen. Cham: Springer International Publishing, 2017, pp. 505–520. ISBN: 978-3-319-57987-0.

[151] *HaptX Gloves G1*. URL: https://haptx.com/. (accessed 24/01/2023).

[152] Daniel Harris, Dominic Potts, and Steven Houben. "User-Elicited Surface and Motion Gestures for Object Manipulation in Mobile Augmented Reality". *Adjunct Publication of the 24th International Conference on Human-Computer Interaction with Mobile Devices and Services*. MobileHCI '22. Vancouver, BC, Canada: Association for Computing Machinery, 2022. ISBN: 9781450393416. DOI: 10.1145/3528575.3551443. URL: https://doi.org/10.1145/3528575.3551443.

[153] Chris Harrison, Shilpa Ramamurthy, and Scott E. Hudson. "On-Body Interaction: Armed and Dangerous". *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*. TEI '12. Kingston, Ontario, Canada: Association for Computing Machinery, 2012, pp. 69–76. ISBN:

9781450311748. DOI: 10.1145/2148131.2148148. URL: https://doi.org/10.1145/2148131.2148148.

[154] Donna Hart and Robert W Sussman. *Man the hunted: Primates, predators, and human evolution*. Routledge, 2018.

[155] Devamardeep Hayatpur, Seongkook Heo, Haijun Xia, Wolfgang Stuerzlinger, and Daniel Wigdor. "Plane, ray, and point: Enabling precise spatial manipulations with shape constraints". *UIST 2019 - Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. 2019. ISBN: 9781450368162. DOI: 10.1145/3332165.3347916.

[156] Steven Henderson and Steven Feiner. "Opportunistic Tangible User Interfaces for Augmented Reality". *IEEE Transactions on Visualization and Computer Graphics* 16.1 (2010), pp. 4–16. DOI: 10.1109/TVCG.2009.91.

[157] Steven J. Henderson and Steven Feiner. "Opportunistic Controls: Leveraging Natural Affordances as Tangible User Interfaces for Augmented Reality". *Proceedings of the 2008 ACM Symposium on Virtual Reality Software and Technology*. VRST '08. Bordeaux, France: Association for Computing Machinery, 2008, pp. 211–218. ISBN: 9781595939517. DOI: 10.1145/1450579.1450625. URL: https://doi.org/10.1145/1450579.1450625.

[158] Kenneth P. Herndon, Andries van Dam, and Michael Gleicher. "The Challenges of 3D Interaction: A CHI '94 Workshop". *SIGCHI Bull.* 26.4 (Oct. 1994), pp. 36–43. ISSN: 0736-6906. DOI: 10.1145/191642.191652. URL: https://doi.org/10.1145/191642.191652.

[159] Anuruddha Hettiarachchi and Daniel Wigdor. "Annexing reality: Enabling opportunistic use of everyday objects as tangible proxies in augmented reality". *Conference on Human Factors in Computing Systems - Proceedings*. 2016. ISBN: 9781450333627. DOI: 10.1145/2858036.2858134.

[160] Mikito Hikosaka and Yu Aramaki. "Neuromuscular Fatigue in Unimanual Handgrip Does Not Completely Affect Simultaneous Bimanual Handgrip". *Frontiers in Human Neuroscience* 15 (2021).

[161] Otmar Hilliges, Dominikus Baur, and Andreas Butz. "Photohelix: Browsing, Sorting and Sharing Digital Photo Collections". *Second Annual IEEE International Workshop on Horizontal Interactive Human-Computer Systems (TABLETOP'07)*. 2007, pp. 87–94. DOI: 10.1109/TABLETOP.2007.14.

[162] Otmar Hilliges, Shahram Izadi, Andrew D. Wilson, Steve Hodges, Armando Garcia-Mendoza, and Andreas Butz. "Interactions in the air: Adding further depth to interactive tabletops". *UIST 2009 - Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology*. 2009. ISBN: 9781605587455. DOI: 10.1145/1622176.1622203.

[163] Otmar Hilliges, David Kim, Shahram Izadi, Malte Weiss, and Andrew Wilson. "HoloDesk: Direct 3d Interactions with a Situated See-through Display". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: Association for Computing Machinery, 2012, pp. 2421–2430. ISBN: 9781450310154. DOI: 10.1145/2207676.2208405. URL: https://doi.org/10.1145/2207676.2208405.

[164] Juan David Hincapié-Ramos, Xiang Guo, Paymahn Moghadasian, and Pourang Irani. "Consumed Endurance: A Metric to Quantify Arm Fatigue of Mid-Air Interactions". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '14. Toronto, Ontario, Canada: Association for Computing Machinery, 2014, pp. 1063–1072. ISBN: 9781450324731. DOI: 10.1145/2556288.2557130. URL: https://doi.org/10.1145/2556288.2557130.

[165] Juan David Hincapié-Ramos, Kasim Ozacar, Pourang P. Irani, and Yoshifumi Kitamura. "GyroWand: IMU-Based Raycasting for Augmented Reality Head-Mounted Displays". *Proceedings of the 3rd ACM Symposium on Spatial User Interaction*. SUI '15. Los Angeles, California, USA: Association for Computing Machinery, 2015, pp. 89–98. ISBN: 9781450337038. DOI: 10.1145/2788940.2788947. URL: https://doi.org/10.1145/2788940.2788947.

[166] Ronan Hinchet, Velko Vechev, Herbert Shea, and Otmar Hilliges. "DextrES: Wearable Haptic Feedback for Grasping in VR via a Thin Form-Factor Electrostatic Brake". *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST '18. Berlin, Germany: Association for Computing Machinery, 2018, pp. 901–912. ISBN: 9781450359481. DOI: 10.1145/3242587.3242657. URL: https://doi.org/10.1145/3242587.3242657.

[167] Ken Hinckley. "Synchronous gestures for multiple persons and computers". *Proceedings of the 16th annual ACM symposium on User interface software and technology*. 2003, pp. 149–158.

[168] Ken Hinckley, Randy Pausch, John C Goble, and Neal F Kassell. "Passive real-world interface props for neurosurgical visualization". *Proceedings of the SIGCHI conference on Human factors in computing systems*. 1994, pp. 452–458.

[169] Ken Hinckley, Randy Pausch, Dennis Proffitt, James Patten, and Neal Kassell. "Cooperative bimanual action". *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*. 1997, pp. 27–34.

[170] Ken Hinckley, Koji Yatani, Michel Pahud, Nicole Coddington, Jenny Rodenhouse, Andy Wilson, Hrvoje Benko, and Bill Buxton. "Pen+ touch= new tools". *Proceedings of the 23nd annual ACM symposium on User interface software and technology*. 2010, pp. 27–36.

[171] Ryuji Hirayama, Diego Martinez Plasencia, Nobuyuki Masuda, and Sriram Subramanian. "A volumetric display for visual, tactile and audio presentation using acoustic trapping". *Nature* 575.7782 (2019), pp. 320–323.

[172] *Hoberman Sphere*. URL: https://www.lakeshorelearning.com/products/science/physical-science/hoberman-spheresup-sup/p/ES130/. (accessed 15/01/2023).

[173] David Holman and Hrvoje Benko. "SketchSpace: Designing Interactive Behaviors with Passive Materials". *CHI '11 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '11. Vancouver, BC, Canada: Association for Computing Machinery, 2011, pp. 1987–1992. ISBN: 9781450302685. DOI: 10.1145/1979742.1979867. URL: https://doi.org/10.1145/1979742.1979867.

[174] Nicholas P Holmes and Charles Spence. "The body schema and multisensory representation (s) of peripersonal space". *Cognitive processing* 5 (2004), pp. 94–105.

[175] Matthias Hoppe, Pascal Knierim, Thomas Kosch, Markus Funk, Lauren Futami, Stefan Schneegass, Niels Henze, Albrecht Schmidt, and Tonja Machulla. "VRHapticDrones: Providing Haptics in Virtual Reality through Quadcopters". *Proceedings of the 17th International Conference on Mobile and Ubiquitous Multimedia*. MUM '18. Cairo, Egypt: Association for Computing Machinery, 2018, pp. 7–18. ISBN: 9781450365949. DOI: 10.1145/3282894.3282898. URL: https://doi.org/10.1145/3282894.3282898.

[176] Michael S. Horn and Robert J. K. Jacob. "Designing Tangible Programming Languages for Classroom Use". *Proceedings of the 1st International Conference on Tangible and Embedded Interaction*. TEI '07. Baton Rouge, Louisiana: Association for Computing Machinery, 2007, pp. 159–162. ISBN: 9781595936196. DOI: 10.1145/1226969.1227003. URL: https://doi.org/10.1145/1226969.1227003.

[177] Kasper Hornbæk and Antti Oulasvirta. "What Is Interaction?" *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. Denver, Colorado, USA: Association for Computing Machinery, 2017, pp. 5040–5052. ISBN: 9781450346559. DOI: 10.1145/3025453.3025765. URL: https://doi.org/10.1145/3025453.3025765.

[178] Eva Hornecker. "Beyond affordance: tangibles' hybrid nature". *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction*. 2012, pp. 175–182.

[179] Steven Houben and Nicolai Marquardt. "WatchConnect: A Toolkit for Prototyping Smartwatch-Centric Cross-Device Applications". *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. Seoul, Republic of Korea: Association for Computing Machinery, 2015, pp. 1247–1256. ISBN: 9781450331456. DOI: 10.1145/2702123.2702215. URL: https://doi.org/10.1145/2702123.2702215.

[180] Elise van den Hoven and Ali Mazalek. "Grasping gestures: Gesturing with physical artifacts". *Artificial Intelligence for Engineering Design, Analysis and Manufacturing: AI EDAM* 25.3 (2011), p. 255.

[181] Elise van den Hoven and Ali Mazalek. "Grasping gestures: Gesturing with physical artifacts". *AI EDAM* 25.3 (2011), pp. 255–271. DOI: 10.1017/S0890060411000072.

[182] Meng-Ju Hsieh, Rong-Hao Liang, Da-Yuan Huang, Jheng-You Ke, and Bing-Yu Chen. "RFIBricks: Interactive Building Blocks Based on RFID". *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: Association for Computing Machinery, 2018, pp. 1–10. ISBN: 9781450356206. DOI: 10.1145/3173574.3173763. URL: https://doi.org/10.1145/3173574.3173763.

[183] Hsuan-Yu Hsueh, Chien-Hua Chen, Irene Chen, Chih-Yuan Yao, and Hung-Kuo Chu. "ARToken: A Tangible Device for Dynamically Binding Real-World Objects with Virtual Representation". *Adjunct Proceedings of the 2021 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2021 ACM International Symposium on Wearable Computers*. UbiComp '21. Virtual, USA: Association for Computing Machinery, 2021, pp. 560–564. ISBN: 9781450384612. DOI: 10.1145/3460418.3480161. URL: https://doi.org/10.1145/3460418.3480161.

[184] Felix Hu, Ariel Zekelman, Michael Horn, and Frances Judd. "Strawbies: Explorations in Tangible Programming". *Proceedings of the 14th International Conference on Interaction Design and Children*. IDC '15. Boston, Massachusetts: Association for Computing Machinery, 2015, pp. 410–413. ISBN: 9781450335904. DOI: 10.1145/2771839.2771866. URL: https://doi.org/10.1145/2771839.2771866.

[185] Edwin L. Hutchins, James D. Hollan, and Donald A. Norman. "Direct Manipulation Interfaces". *Human–Computer Interaction* 1.4 (1985), pp. 311–338. DOI: 10.1207/s15327051hci0104_2. eprint: https://doi.org/10.1207/s15327051hci0104_2. URL: https://doi.org/10.1207/s15327051hci0104_2.

[186] Hilary Hutchinson, Wendy Mackay, Bo Westerlund, Benjamin B. Bederson, Allison Druin, Catherine Plaisant, Michel Beaudouin-Lafon, Stéphane Conversy, Helen Evans, Heiko Hansen, Nicolas Roussel, and Björn Eiderbäck. "Technology Probes: Inspiring Design for and with Families". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '03. Ft. Lauderdale, Florida, USA: Association for Computing Machinery, 2003, pp. 17–24. ISBN: 1581136307. DOI: 10.1145/642611.642616. URL: https://doi.org/10.1145/642611.642616.

[187] *Infinity Cube*. URL: https://www.amazon.co.uk/Funxim-Infinity-suitable-Version-Infinite/dp/B078NNQFLV?ref_=Oct_d_otopr_d_26894543031_2&pd_rd_w=jLuTl&content-id=amzn1.sym.24a10ac9-2af8-4c2b-90b8-d76667752374&pf_rd_p=24a10ac9-2af8-4c2b-90b8-d76667752374&pf_rd_r=M8XP6NFN5PW9HCZ2SGNR&pd_rd_wg=faWbO&pd_rd_r=16cb5a54-50d1-4248-9af3-c7b7123454eb&pd_rd_i=B078NNQFLV. (accessed 15/01/2023).

[188] Computer Dictionary of Information Technology. *Gorilla Arm*. URL: https://www.computer-dictionary-online.org/definitions-g/gorilla-arm.html. (accessed: 30/11/2022).

[189] Hiroshi Ishii. "Tangible bits: beyond pixels". *Proceedings of the 2nd international conference on Tangible and embedded interaction*. 2008, pp. xv–xxv.

[190] Hiroshi Ishii, Minoru Kobayashi, and Kazuho Arita. "Iterative design of seamless collaboration media". *Communications of the ACM* 37.8 (1994), pp. 83–97.

[191] Hiroshi Ishii and Brygg Ullmer. "Tangible bits: Towards seamless interfaces between people, bits and atoms". *Conference on Human Factors in Computing Systems - Proceedings*. 1997.

[192] Paul Issartel, Lonni Besançon, Tobias Isenberg, and Mehdi Ammi. "A Tangible Volume for Portable 3D Interaction". *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)*. 2016, pp. 215–220. DOI: 10.1109/ISMAR-Adjunct.2016.0079.

[193] Bret Jackson and Daniel F. Keefe. "Lift-Off: Using Reference Imagery and Freehand Sketching to Create 3D Models in VR". *IEEE Transactions on Visualization and Computer Graphics* 22.4 (2016), pp. 1442–1451. DOI: 10.1109/TVCG.2016.2518099.

[194] Robert J.K. Jacob, Audrey Girouard, Leanne M. Hirshfield, Michael S. Horn, Orit Shaer, Erin Treacy Solovey, and Jamie Zigelbaum. "Reality-Based Interaction: A Framework for Post-WIMP Interfaces". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '08. Florence, Italy: Association for Computing Machinery, 2008, pp. 201–210. ISBN: 9781605580111. DOI: 10.1145/1357054.1357089. URL: https://doi.org/10.1145/1357054.1357089.

[195] Sujin Jang, Wolfgang Stuerzlinger, Satyajit Ambike, and Karthik Ramani. "Modeling Cumulative Arm Fatigue in Mid-Air Interaction Based on Perceived Exertion and Kinetics of Arm Motion". *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. Denver, Colorado, USA: Association for Computing Machinery, 2017, pp. 3328–3339. ISBN: 9781450346559. DOI: 10.1145/3025453.3025523. URL: https://doi.org/10.1145/3025453.3025523.

[196] Sung-A Jang, Graham Wakefield, and Sung-Hee Lee. "Incorporating Kinesthetic Creativity and Gestural Play into Immersive Modeling". *Proceedings of the 4th International Conference on Movement Computing*. MOCO '17. London, United Kingdom: Association for Computing Machinery, 2017. ISBN: 9781450352093. DOI: 10.1145/3077981.3078045. URL: https://doi.org/10.1145/3077981.3078045.

[197] Ying Jiang, Congyi Zhang, Hongbo Fu, Alberto Cannavò, Fabrizio Lamberti, Henry Y K Lau, and Wenping Wang. "HandPainter - 3D Sketching in VR with Hand-Based Physical Proxy". *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. Yokohama, Japan: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445302. URL: https://doi.org/10.1145/3411764.3445302.

[198] Kenneth O Johnson. "The roles and functions of cutaneous mechanoreceptors". *Current Opinion in Neurobiology* 11.4 (2001), pp. 455–461. ISSN: 0959-4388. DOI: https://doi.org/10.1016/S0959-4388(00)00234-8. URL: https://www.sciencedirect.com/science/article/pii/S0959438800002348.

[199] Sergi Jordà, Günter Geiger, Marcos Alonso, and Martin Kaltenbrunner. "The reacTable: exploring the synergy between live music performance and tabletop

tangible interfaces". *Proceedings of the 1st international conference on Tangible and embedded interaction*. 2007, pp. 139–146.

[200] Ed Kaiser, Alex Olwal, David McGee, Hrvoje Benko, Andrea Corradini, Xiaoguang Li, Phil Cohen, and Steven Feiner. "Mutual Disambiguation of 3D Multimodal Interaction in Augmented and Virtual Reality". *Proceedings of the 5th International Conference on Multimodal Interfaces*. ICMI '03. Vancouver, British Columbia, Canada: Association for Computing Machinery, 2003, pp. 12–19. ISBN: 1581136218. DOI: 10.1145/958432.958438. URL: https://doi.org/10.1145/958432.958438.

[201] Martin Kaltenbrunner and Ross Bencina. "ReacTIVision: A Computer-Vision Framework for Table-Based Tangible Interaction". *Proceedings of the 1st International Conference on Tangible and Embedded Interaction*. TEI '07. Baton Rouge, Louisiana: Association for Computing Machinery, 2007, pp. 69–74. ISBN: 9781595936196. DOI: 10.1145/1226969.1226983. URL: https://doi.org/10.1145/1226969.1226983.

[202] Maria Karam and m. c. schraefel. *A Taxonomy of Gestures in Human Computer Interactions*. Project Report. 2005. URL: https://eprints.soton.ac.uk/261149/.

[203] H. Kato and M. Billinghurst. "Marker tracking and HMD calibration for a video-based augmented reality conferencing system". *Proceedings 2nd IEEE and ACM International Workshop on Augmented Reality (IWAR'99)*. 1999, pp. 85–94. DOI: 10.1109/IWAR.1999.803809.

[204] H. Kato, M. Billinghurst, I. Poupyrev, K. Imamoto, and K. Tachibana. "Virtual object manipulation on a table-top AR environment". *Proceedings IEEE and ACM International Symposium on Augmented Reality (ISAR 2000)*. 2000, pp. 111–119. DOI: 10.1109/ISAR.2000.880934.

[205] Oliver Beren Kaul and Michael Rohs. "HapticHead: A Spherical Vibrotactile Grid around the Head for 3D Guidance in Virtual and Augmented Reality". *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. Denver, Colorado, USA: Association for Computing Machinery, 2017, pp. 3729–3740. ISBN: 9781450346559. DOI: 10.1145/3025453.3025684. URL: https://doi.org/10.1145/3025453.3025684.

[206] Annie Kelly, R. Benjamin Shapiro, Jonathan de Halleux, and Thomas Ball. "ARcadia: A Rapid Prototyping Platform for Real-Time Tangible Interfaces". *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: Association for Computing Machinery, 2018, pp. 1–8. ISBN: 9781450356206. DOI: 10.1145/3173574.3173983. URL: https://doi.org/10.1145/3173574.3173983.

[207] Hojoong Kim, Young-Tae Kwon, Hyo-Ryoung Lim, Jong-Hoon Kim, Yun-Soung Kim, and Woon-Hong Yeo. "Recent advances in wearable sensors and integrated functional devices for virtual and augmented reality applications". *Advanced Functional Materials* 31.39 (2021), p. 2005692.

[208] Jonghyun Kim, Youngmo Jeong, Michael Stengel, Kaan Aksit, Rachel A Albert, Ben Boudaoud, Trey Greer, Joohwan Kim, Ward Lopes, Zander Majercik, et al. "Foveated AR: dynamically-foveated augmented reality display." *ACM Trans. Graph.* 38.4 (2019), pp. 99–1.

[209] Kangsoo Kim, Mark Billinghurst, Gerd Bruder, Henry Been-Lirn Duh, and Gregory F. Welch. "Revisiting Trends in Augmented Reality Research: A Review of the 2nd Decade of ISMAR (2008–2017)". *IEEE Transactions on Visualization and Computer Graphics* 24.11 (2018), pp. 2947–2962. DOI: 10.1109/TVCG.2018.2868591.

[210] Mi Jeong Kim and Mary Lou Maher. "The impact of tangible user interfaces on designers' spatial cognition". *Human-Computer Interaction* (2008). ISSN: 07370024. DOI: 10.1080/07370020802016415.

[211] MyoungGon Kim and JungHyun Han. "Effects of Switchable DOF for Mid-Air Manipulation in Immersive Virtual Environments". *International Journal of Human–Computer Interaction* 35.13 (2019), pp. 1147–1159. DOI: 10.1080/10447318.2018.1514163. eprint: https://doi.org/10.1080/10447318.2018.1514163. URL: https://doi.org/10.1080/10447318.2018.1514163.

[212] David Kirsh. "Embodied cognition and the magical future of interaction design". *ACM Transactions on Computer-Human Interaction (TOCHI)* 20.1 (2013), pp. 1–30.

[213] Pascal Knierim, Thomas Kosch, and Albrecht Schmidt. "The Nomadic Office: A Location Independent Workspace Through Mixed Reality". *IEEE Pervasive Computing* 20.4 (2021), pp. 71–78. DOI: 10.1109/MPRV.2021.3119378.

[214] Yukari Konishi, Nobuhisa Hanamitsu, Benjamin Outram, Kouta Minamizawa, Tetsuya Mizuguchi, and Ayahiko Sato. "Synesthesia Suit: The Full Body Immersive Experience". *ACM SIGGRAPH 2016 VR Village*. SIGGRAPH '16. Anaheim, California: Association for Computing Machinery, 2016. ISBN: 9781450343770. DOI: 10.1145/2929490.2932629. URL: https://doi.org/10.1145/2929490.2932629.

[215] Christian Kray, Daniel Nesbitt, John Dawson, and Michael Rohs. "User-defined gestures for connecting mobile phones, public displays, and tabletops". *ACM International Conference Proceeding Series*. 2010. ISBN: 9781605588353. DOI: 10.1145/1851600.1851640.

[216] Max Krichenbauer, Goshiro Yamamoto, Takafumi Taketom, Christian Sandor, and Hirokazu Kato. "Augmented Reality versus Virtual Reality for 3D Object Manipulation". *IEEE Transactions on Visualization and Computer Graphics* 24.2 (2018), pp. 1038–1048. DOI: 10.1109/TVCG.2017.2658570.

[217] Myron W. Krueger, Thomas Gionfriddo, and Katrin Hinrichsen. "VIDEOPLACE—an Artificial Reality". *SIGCHI Bull.* 16.4 (Apr. 1985), pp. 35–40. ISSN: 0736-6906. DOI: 10.1145/1165385.317463. URL: https://doi.org/10.1145/1165385.317463.

[218] Kari Kuutti et al. "Activity theory as a potential framework for human-computer interaction research". *Context and consciousness: Activity theory and human-computer interaction* 1744 (1996).

[219] Mikko Kytö, Barrett Ens, Thammathip Piumsomboon, Gun A. Lee, and Mark Billinghurst. "Pinpointing: Precise Head- and Eye-Based Target Selection for Augmented Reality". *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: Association for Computing Machinery, 2018, pp. 1–14. ISBN: 9781450356206. DOI: 10.1145/3173574.3173655. URL: https://doi.org/10.1145/3173574.3173655.

[220] Thomas Langerak, Juan José Zárate, David Lindlbauer, Christian Holz, and Otmar Hilliges. "Omni: Volumetric Sensing and Actuation of Passive Magnetic Tools for Dynamic Haptic Feedback". *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. 2020, pp. 594–606.

[221] Joseph J. LaViola, Ernst Kruijff, Ryan P. McMahan, Doug Bowman, and Ivan P. Poupyrev. *3D User Interfaces: Theory and Practice*. en. 2nd edition. Addison-Wesley usability and HCI series. Boston: Addison-Wesley, 2017, p. 624. ISBN: 978-0-13-403432-4.

[222] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. *Research methods in human-computer interaction*. Morgan Kaufmann, 2017.

[223] Mathieu Le Goc, Lawrence H Kim, Ali Parsaei, Jean-Daniel Fekete, Pierre Dragicevic, and Sean Follmer. "Zooids: Building blocks for swarm user interfaces". *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 2016, pp. 97–109.

[224] Mathieu Le Goc, Allen Zhao, Ye Wang, Griffin Dietz, Rob Semmens, and Sean Follmer. "Investigating Active Tangibles and Augmented Reality for Creativity Support in Remote Collaboration". *Design Thinking Research*. Springer, 2020, pp. 185–200.

[225] Susan J Lederman and Roberta L Klatzky. "Hand movements: A window into haptic object recognition". *Cognitive psychology* 19.3 (1987), pp. 342–368.

[226] David Ledo, Steven Houben, Jo Vermeulen, Nicolai Marquardt, Lora Oehlberg, and Saul Greenberg. "Evaluation Strategies for HCI Toolkit Research". *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: Association for Computing Machinery, 2018, pp. 1–17. ISBN: 9781450356206. DOI: 10.1145/3173574.3173610. URL: https://doi.org/10.1145/3173574.3173610.

[227] Byungjoo Lee, Qiao Deng, Eve Hoggan, and Antti Oulasvirta. "Boxer: A Multimodal Collision Technique for Virtual Objects". *Proceedings of the 19th ACM International Conference on Multimodal Interaction*. ICMI '17. Glasgow, UK: Association for Computing Machinery, 2017, pp. 252–260. ISBN: 9781450355438. DOI: 10.1145/3136755.3136761. URL: https://doi.org/10.1145/3136755.3136761.

[228] Hyeongmook Lee, Mark Billinghurst, and Woontack Woo. "Two-handed tangible interaction techniques for composing augmented blocks". *Virtual Reality* 15.2-3 (2011), pp. 133–146.

[229] Hyun-Jean Lee, Hyungsin Kim, Gaurav Gupta, and Ali Mazalek. "WiiArts: creating collaborative art experience with WiiRemote interaction". *Proceedings of the 2nd international conference on Tangible and embedded interaction*. 2008, pp. 33–36.

[230] Jae Yeol Lee, Gue Won Rhee, and Dong Woo Seo. "Hand gesture-based tangible interactions for manipulating virtual objects in a mixed reality environment". *International Journal of Advanced Manufacturing Technology* (2010). ISSN: 02683768. DOI: 10.1007/s00170-010-2671-x.

[231] Jae Yeol Lee, Gue Won Rhee, and Dong Woo Seo. "Hand gesture-based tangible interactions for manipulating virtual objects in a mixed reality environment". *The International Journal of Advanced Manufacturing Technology* 51.9 (2010), pp. 1069–1082.

[232] Kangdon Lee. "Augmented reality in education and training". *TechTrends* 56.2 (2012), pp. 13–21.

[233] M. Lee, R. Green, and M. Billinghurst. "3D natural hand interaction for AR applications". *2008 23rd International Conference Image and Vision Computing New Zealand*. 2008, pp. 1–6. DOI: 10.1109/IVCNZ.2008.4762125.

[234] Taehee Lee and Tobias Hollerer. "Handy AR: Markerless Inspection of Augmented Reality Objects Using Fingertip Tracking". *2007 11th IEEE International Symposium on Wearable Computers*. 2007, pp. 83–90. DOI: 10.1109/ISWC.2007.4373785.

[235] Kevin Lefeuvre, Soeren Totzauer, Michael Storz, Albrecht Kurze, Andreas Bischof, and Arne Berger. "Bricks, Blocks, Boxes, Cubes, and Dice: On the Role of Cubic Shapes for the Design of Tangible Interactive Devices". *Proceedings of the 2018 Designing Interactive Systems Conference*. DIS '18. Hong Kong, China: Association for Computing Machinery, 2018, pp. 485–496. ISBN: 9781450351980. DOI: 10.1145/3196709.3196768. URL: https://doi.org/10.1145/3196709.3196768.

[236] Daniel Leithinger and Hiroshi Ishii. "Relief: A Scalable Actuated Shape Display". *Proceedings of the Fourth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '10. Cambridge, Massachusetts, USA: Association for Computing Machinery, 2010, pp. 221–222. ISBN: 9781605588414. DOI: 10.1145/1709886.1709928. URL: https://doi.org/10.1145/1709886.1709928.

[237] Daniele Leonardis, Massimiliano Solazzi, Ilaria Bortone, and Antonio Frisoli. "A 3-RSR Haptic Wearable Device for Rendering Fingertip Contact Forces". *IEEE Transactions on Haptics* 10.3 (2017), pp. 305–316. DOI: 10.1109/TOH.2016.2640291.

[238] R.W. Lindeman, J.L. Sibert, and J.K. Hahn. "Hand-held windows: towards effective 2D interaction in immersive virtual environments". *Proceedings IEEE Virtual Reality (Cat. No. 99CB36316)*. 1999, pp. 205–212. DOI: 10.1109/VR.1999.756952.

[239] David Lindlbauer, Jens Emil Grønbæk, Morten Birk, Kim Halskov, Marc Alexa, and Jörg Müller. "Combining Shape-Changing Interfaces and Spatial Augmented Reality Enables Extended Object Appearance". *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. San Jose, California, USA: Association for Computing Machinery, 2016, pp. 791–802. ISBN: 9781450333627. DOI: 10.1145/2858036.2858457. URL: https://doi.org/10.1145/2858036.2858457.

[240] Jo-Yu Lo, Da-Yuan Huang, Chen-Kuo Sun, Chu-En Hou, and Bing-Yu Chen. "RollingStone: Using Single Slip Taxel for Enhancing Active Finger Exploration with a Virtual Reality Controller". *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST '18. Berlin, Germany: Association for Computing Machinery, 2018, pp. 839–851. ISBN: 9781450359481. DOI: 10.1145/3242587.3242627. URL: https://doi.org/10.1145/3242587.3242627.

[241] Benjamin Long, Sue Ann Seah, Tom Carter, and Sriram Subramanian. "Rendering Volumetric Haptic Shapes in Mid-Air Using Ultrasound". *ACM*

Trans. Graph. 33.6 (Nov. 2014). ISSN: 0730-0301. DOI: 10.1145/2661229.2661257. URL: https://doi.org/10.1145/2661229.2661257.

[242] Julian Looser, Mark Billinghurst, Raphaël Grasset, and Andy Cockburn. "An Evaluation of Virtual Lenses for Object Selection in Augmented Reality". *Proceedings of the 5th International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia*. GRAPHITE '07. Perth, Australia: Association for Computing Machinery, 2007, pp. 203–210. ISBN: 9781595939128. DOI: 10.1145/1321261.1321297. URL: https://doi.org/10.1145/1321261.1321297.

[243] Pedro Lopes, Patrik Jonell, and Patrick Baudisch. "Affordance++ : allowing objects to communicate dynamic use". *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*. 2015. ISBN: 9781450331456. DOI: 10.1145/2702123.2702128. arXiv: arXiv:1011.1669v3.

[244] Pedro Lopes, Sijing You, Lung Pan Cheng, Sebastian Marwecki, and Patrick Baudisch. "Providing haptics to walls & heavy objects in virtual reality by means of electrical muscle stimulation". *Conference on Human Factors in Computing Systems - Proceedings*. 2017. ISBN: 9781450346559. DOI: 10.1145/3025453.3025600.

[245] I. Scott MacKenzie. "Fitts' Law as a Research and Design Tool in Human-Computer Interaction". *Human-Computer Interaction* 7.1 (1992), pp. 91–139. DOI: 10.1207/s15327051hci0701_3. eprint: https://doi.org/10.1207/s15327051hci0701_3. URL: https://doi.org/10.1207/s15327051hci0701_3.

[246] *Magic Leap 2*. URL: https://www.magicleap.com/magic-leap-2. (accessed: 01.12.2022).

[247] *Magic Leap: One*. URL: https://www.magicleap.com/en-us/magic-leap-1. (accessed 15/01/2023).

[248] *Magnetic Balls*. URL: https://www.manomano.co.uk/catalogue/p/1000pcs-3mm-magnetic-ball-set-magic-magnet-cube-building-toy-for-stress-relief-mix-10-color-for-10-years-old-64184083. (accessed 15/01/2023).

[249] *Magnetic Cubes*. URL: https://magnetstore.co.uk/product/5mm-x-5mm-x-5mm-cube-neodymium-magnet/. (accessed 15/01/2023).

[250] *Magnetic Spinner*. URL: https://www.aliexpress.com/item/1005005026546426.html?spm=a2g0o.productlist.main.9.3901mSlumSluUa&algo_pvid=791eaf60-e67b-4788-93d2-54e94938d1a5&algo_exp_id=791eaf60-e67b-4788-93d2-54e94938d1a5-4&pdp_ext_f=%5C%7B%5C%22sku_id%5C%22%5C%3A%5C%2212000031381277605%5C%22%5C%7D&pdp_npi=3%5C%40dis%5C%21GBP%5C%216.14%5C%213.99%5C%21%5C%21%5C%21%5C%21%5C%21%5C%402100b69816762973300641118d06d2%5C%2112000031381277605%5C%21sea%5C%21UK%5C%21759993656&curPageLogUid=spumh59XoiV4. (accessed 15/01/2023).

[251] Mary Lou Maher and Lina Lee. "Designing for gesture and tangible interaction". *Synthesis Lectures on Human-Centered Interaction* 10.2 (2017), pp. i–111.

[252] S. Malik, C. McDonald, and G. Roth. "Hand tracking for interactive pattern-based augmented reality". *Proceedings. International Symposium on Mixed and Augmented Reality*. 2002, pp. 117–126. DOI: 10.1109/ISMAR.2002.1115080.

[253] Federico Manuri and Andrea Sanna. "A survey on applications of augmented reality". *ACSIJ Advances in Computer Science: an International Journal* 5.1 (2016), pp. 18–27.

[254] *Marble in Mesh*. URL: https://www.multi-sensoryworld.co.uk/products/marble-fidget. (accessed 15/01/2023).

[255] Diako Mardanbegi, Benedikt Mayer, Ken Pfeuffer, Shahram Jalaliniya, Hans Gellersen, and Alexander Perzl. "EyeSeeThrough: Unifying Tool Selection and Application in Virtual Environments". *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 2019, pp. 474–483. DOI: 10.1109/VR.2019.8797988.

[256] Nicolai Marquardt, Robert Diaz-Marino, Sebastian Boring, and Saul Greenberg. "The Proximity Toolkit: Prototyping Proxemic Interactions in Ubiquitous Computing Ecologies". *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*. UIST '11. Santa Barbara, California, USA: Association for Computing Machinery, 2011, pp. 315–326. ISBN: 9781450307161. DOI: 10.1145/2047196.2047238. URL: https://doi.org/10.1145/2047196.2047238.

[257] Riccardo Masoni, Francesco Ferrise, Monica Bordegoni, Michele Gattullo, Antonio E. Uva, Michele Fiorentino, Ernesto Carrabba, and Michele Di Donato. "Supporting Remote Maintenance in Industry 4.0 through Augmented Reality". *Procedia Manufacturing* 11 (2017). 27th International Conference on Flexible Automation and Intelligent Manufacturing, FAIM2017, 27-30 June 2017, Modena, Italy, pp. 1296–1302. ISSN: 2351-9789. DOI: https://doi.org/

10.1016/j.promfg.2017.07.257. URL: https://www.sciencedirect.com/science/article/pii/S2351978917304651.

[258] Atsushi Matsubayashi, Yasutoshi Makino, and Hiroyuki Shinoda. "Direct Finger Manipulation of 3D Object Image with Ultrasound Haptic Feedback". *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. Glasgow, Scotland Uk: Association for Computing Machinery, 2019, pp. 1–11. ISBN: 9781450359702. DOI: 10.1145/3290605.3300317. URL: https://doi.org/10.1145/3290605.3300317.

[259] *Maya by Autodesk*. URL: https://www.autodesk.co.uk/products/maya/overview. (accessed 20/01/2023).

[260] John C. McClelland, Robert J. Teather, and Audrey Girouard. "Haptobend: Shape-Changing Passive Haptic Feedback in Virtual Reality". *Proceedings of the 5th Symposium on Spatial User Interaction*. SUI '17. Brighton, United Kingdom: Association for Computing Machinery, 2017, pp. 82–90. ISBN: 9781450354868. DOI: 10.1145/3131277.3132179. URL: https://doi.org/10.1145/3131277.3132179.

[261] Tim McGraw, Esteban Garcia, and Drew Sumner. "Interactive Swept Surface Modeling in Virtual Reality with Motion-Tracked Controllers". *Proceedings of the Symposium on Sketch-Based Interfaces and Modeling*. SBIM '17. Los Angeles, California: Association for Computing Machinery, 2017. ISBN: 9781450350792. DOI: 10.1145/3092907.3092908. URL: https://doi.org/10.1145/3092907.3092908.

[262] Jess McIntosh, Hubert Dariusz Zajac, Andreea Nicoleta Stefan, Joanna Bergström, and Kasper Hornbæk. "Iteratively Adapting Avatars Using Task-Integrated Optimisation". *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. UIST '20. Virtual Event, USA: Association for Computing Machinery, 2020, pp. 709–721. ISBN: 9781450375146. DOI: 10.1145/3379337.3415832. URL: https://doi.org/10.1145/3379337.3415832.

[263] Gianluca Memoli, Mihai Caleap, Michihiro Asakawa, Deepak R Sahoo, Bruce W Drinkwater, and Sriram Subramanian. "Metamaterial bricks and quantization of meta-surfaces". *Nature communications* 8.1 (2017), pp. 1–8.

[264] D. Mendes, F. M. Caputo, A. Giachetti, A. Ferreira, and J. Jorge. "A Survey on 3D Virtual Object Manipulation: From the Desktop to Immersive Virtual Environments". *Computer Graphics Forum* 38.1 (2019), pp. 21–45. DOI: https://doi.org/10.1111/cgf.13390. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13390. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13390.

[265] Daniel Mendes, Filipe Relvas, Alfredo Ferreira, and Joaquim Jorge. "The Benefits of DOF Separation in Mid-Air 3D Object Manipulation". *Proceedings of the 22nd ACM Conference on Virtual Reality Software and Technology*. VRST '16. Munich, Germany: Association for Computing Machinery, 2016, pp. 261–268. ISBN: 9781450344913. DOI: 10.1145/2993369.2993396. URL: https://doi.org/10.1145/2993369.2993396.

[266] *Microsoft HoloLens*. URL: https://www.microsoft.com/en-us/hololens. (accessed 15/01/2023).

[267] *Microsoft HoloLens: Partner Demo with Maya by Autodesk*. 2015. URL: https://www.youtube.com/watch?v=pdh-9pOlwpQ. (accessed 20/01/2023).

[268] Paul E Mieschke, Digby Elliott, Werner F Helsen, Richard G Carson, and Jamie A Coull. "Manual asymmetries in the preparation and control of goal-directed movements". *Brain and cognition* 45.1 (2001), pp. 129–140.

[269] Paul Milgram and Fumio Kishino. "A taxonomy of mixed reality visual displays". *IEICE TRANSACTIONS on Information and Systems* 77.12 (1994), pp. 1321–1329.

[270] Paul Milgram, Haruo Takemura, Akira Utsumi, and Fumio Kishino. "Augmented reality: a class of displays on the reality-virtuality continuum". *Telemanipulator and Telepresence Technologies*. Ed. by Hari Das. Vol. 2351. International Society for Optics and Photonics. SPIE, 1995, pp. 282–292. DOI: 10.1117/12.197321. URL: https://doi.org/10.1117/12.197321.

[271] Alexandre Millette and Michael J. McGuffin. "DualCAD: Integrating Augmented Reality with a Desktop GUI and Smartphone Interaction". *2016 IEEE International Symposium on Mixed and Augmented Reality (ISMAR-Adjunct)*. 2016, pp. 21–26. DOI: 10.1109/ISMAR-Adjunct.2016.0030.

[272] Maria L. Montoya Freire, Dominic Potts, Niraj Ramesh Dayama, Antti Oulasvirta, and Mario Di Francesco. "Foraging-based optimization of pervasive displays". *Pervasive and Mobile Computing* 55 (2019), pp. 45–58. ISSN: 1574-1192. DOI: https://doi.org/10.1016/j.pmcj.2019.02.008. URL: https://www.sciencedirect.com/science/article/pii/S1574119218307570.

[273] Thomas Muender, Anke V. Reinschluessel, Sean Drewes, Dirk Wenig, Tanja Döring, and Rainer Malaka. "Does it feel real? Using tangibles with different fidelities to build and explore scenes in virtual reality". *Conference on Human Factors in Computing Systems - Proceedings*. 2019. ISBN: 9781450359702. DOI: 10.1145/3290605.3300903.

[274] *Myron Krueger's VIDEOPLACE Demo*. URL: https://aboutmyronkrueger.weebly.com/videoplace.html. (accessed 20/01/2023).

[275] John Napier, John Russell Napier, and Russell H Tuttle. *Hands*. Princeton University Press, 1993.

[276] Alexander Nelson, Gurashish Singh, Ryan Robucci, Chintan Patel, and Nilanjan Banerjee. "Adaptive and Personalized Gesture Recognition Using Textile Capacitive Sensor Arrays". *IEEE Transactions on Multi-Scale Computing Systems* 1.2 (2015), pp. 62–75. DOI: 10.1109/TMSCS.2015.2495100.

[277] Adrian K. T. Ng, Leith K. Y. Chan, and Henry Y. K. Lau. "A low-cost lighthouse-based virtual reality head tracking system". *2017 International Conference on 3D Immersion (IC3D)*. 2017, pp. 1–5. DOI: 10.1109/IC3D.2017.8251910.

[278] Jakob Nielsen and Rolf Molich. "Heuristic Evaluation of User Interfaces". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '90. Seattle, Washington, USA: Association for Computing Machinery, 1990, pp. 249–256. ISBN: 0201509326. DOI: 10.1145/97243.97281. URL: https://doi.org/10.1145/97243.97281.

[279] SS Muhammad Nizam, Rimaniza Zainal Abidin, Nurhazarifah Che Hashim, Meng Chun Lam, Haslina Arshad, and NAA Majid. "A review of multimodal interaction technique in augmented reality environment". *Int. J. Adv. Sci. Eng. Inf. Technol* 8.4-2 (2018), p. 1460.

[280] Mojtaba Noghabaei, Arsalan Heydarian, Vahid Balali, and Kevin Han. "Trend Analysis on Adoption of Virtual and Augmented Reality in the Architecture, Engineering, and Construction Industry". *Data* 5.1 (2020). ISSN: 2306-5729. DOI: 10.3390/data5010026. URL: https://www.mdpi.com/2306-5729/5/1/26.

[281] Don Norman. *The design of everyday things: Revised and expanded edition*. Basic books, 2013.

[282] Don Norman. *Things that make us smart: Defending human attributes in the age of the machine*. Diversion Books, 2014.

[283] Donald A Norman. *The psychology of everyday things*. Basic books, 1988.

[284] Donald A Norman. "Cognitive artifacts". *Designing interaction: Psychology at the human-computer interface* 1.1 (1991), pp. 17–38.

[285] Donald A Norman and Stephen W Draper. "User centered system design: New perspectives on human-computer interaction" (1986).

[286] Diana Nowacka, Karim Ladha, Nils Y Hammerla, Daniel Jackson, Cassim Ladha, Enrico Rukzio, and Patrick Olivier. "Touchbugs: Actuated tangibles on multi-touch tables". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2013, pp. 759–762.

[287] Benjamin Nuernberger, Eyal Ofek, Hrvoje Benko, and Andrew D. Wilson. "SnapToReality: Aligning Augmented Reality to the Real World". *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. San Jose, California, USA: Association for Computing Machinery, 2016, pp. 1233–1244. ISBN: 9781450333627. DOI: 10.1145/2858036.2858250. URL: https://doi.org/10.1145/2858036.2858250.

[288] Ji-Young Oh and Wolfgang Stuerzlinger. "Moving objects with 2D input devices in CAD systems and desktop virtual environments". *Proceedings of Graphics Interface 2005*. 2005, pp. 195–202.

[289] Simon Olberding, Nan-Wei Gong, John Tiab, Joseph A. Paradiso, and Jürgen Steimle. "A Cuttable Multi-Touch Sensor". *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*. UIST '13. St. Andrews, Scotland, United Kingdom: Association for Computing Machinery, 2013, pp. 245–254. ISBN: 9781450322683. DOI: 10.1145/2501988.2502048. URL: https://doi.org/10.1145/2501988.2502048.

[290] Bjørnar Olsen. *In defense of things: archaeology and the ontology of objects*. Rowman Altamira, 2010.

[291] A. Olwal, H. Benko, and S. Feiner. "SenseShapes: using statistical geometry for object selection in a multimodal augmented reality". *The Second IEEE and ACM International Symposium on Mixed and Augmented Reality, 2003. Proceedings*. 2003, pp. 300–301. DOI: 10.1109/ISMAR.2003.1240730.

[292] Alex Olwal and Steven Feiner. "The Flexible Pointer: An Interaction Technique for Selection in Augmented and Virtual Reality". *Current Science* (2011). ISSN: 00113891. DOI: 10.1.1.214.629.

[293] Alex Olwal, Jonny Gustafsson, and Christoffer Lindfors. "Spatial augmented reality on industrial CNC-machines". *The Engineering Reality of Virtual Reality 2008*. Ed. by Ian E. McDowall and Margaret Dolinsky. Vol. 6804. International Society for Optics and Photonics. SPIE, 2008, p. 680409. DOI: 10.1117/12.760960. URL: https://doi.org/10.1117/12.760960.

[294] *ONO Roller*. URL: https://www.toyscentral.uk/products/ono-roller-grey-the-original-handheld-fidget-toy-for-adults-help-relieve-stress-anxiety-tension-promotes-focus-clarity-compact-portable-design. (accessed 15/01/2023).

[295] Rachel Ormston, Liz Spencer, Matt Barnard, and Dawn Snape. "The foundations of qualitative research". *Qualitative research practice: A guide for social science students and researchers* 2 (2014), pp. 52–55.

[296] Francisco R. Ortega, Katherine Tarre, Armando B. Barreto, Mathew Kress, Naphtali D. Rishe, and Adam S. Williams. "Selection and manipulation whole body gesture elicitation study in virtual reality". *26th IEEE Conference on Virtual Reality and 3D User Interfaces, VR 2019 - Proceedings* (2019), pp. 1110–1111. DOI: 10.1109/VR.2019.8798182.

[297] Sergio Orts-Escolano et al. "Holoportation: Virtual 3D Teleportation in Real-Time". *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. UIST '16. Tokyo, Japan: Association for Computing Machinery, 2016, pp. 741–754. ISBN: 9781450341899. DOI: 10.1145/2984511.2984517. URL: https://doi.org/10.1145/2984511.2984517.

[298] Antti Oulasvirta and Kasper Hornbæk. "HCI Research as Problem-Solving". *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. CHI '16. San Jose, California, USA: Association for Computing Machinery, 2016, pp. 4956–4967. ISBN: 9781450333627. DOI: 10.1145/2858036.2858283. URL: https://doi.org/10.1145/2858036.2858283.

[299] Volker Paelke. "Augmented reality in the smart factory: Supporting workers in an industry 4.0. environment". *Proceedings of the 2014 IEEE Emerging Technology and Factory Automation (ETFA)*. 2014, pp. 1–4. DOI: 10.1109/ETFA.2014.7005252.

[300] *Paper*. URL: https://www.amazon.co.uk/100-Sheets-160gsm-White-Card/dp/B00TQUPXKM. (accessed 15/01/2023).

[301] Youngmin Park, Vincent Lepetit, and Woontack Woo. "Multiple 3D Object tracking for augmented reality". *2008 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*. 2008, pp. 117–120. DOI: 10.1109/ISMAR.2008.4637336.

[302] Esben Warming Pedersen and Kasper Hornbæk. "Tangible Bots: Interaction with Active Tangibles in Tabletop Interfaces". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: Association for Computing Machinery, 2011, pp. 2975–2984. ISBN: 9781450302289. DOI: 10.1145/1978942.1979384. URL: https://doi.org/10.1145/1978942.1979384.

[303] Thomas Pederson, Lars-Erik Janlert, and Dipak Surie. "a Situative Space model for mobile mixed-reality Computing". *IEEE pervasive computing* 10.4 (2010), pp. 73–83.

[304] Siyou Pei, Alexander Chen, Jaewook Lee, and Yang Zhang. "Hand Interfaces: Using Hands to Imitate Objects in AR/VR for Expressive Interactions". *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI '22. New Orleans, LA, USA: Association for Computing Machinery, 2022. ISBN: 9781450391573. DOI: 10.1145/3491102.3501898. URL: https://doi.org/10.1145/3491102.3501898.

[305] J. Perret and E. Vander Poorten. "Touching Virtual Reality: A Review of Haptic Gloves". *ACTUATOR 2018; 16th International Conference on New Actuators*. 2018, pp. 1–5.

[306] Ken Pfeuffer, Benedikt Mayer, Diako Mardanbegi, and Hans Gellersen. "Gaze + Pinch Interaction in Virtual Reality". *Proceedings of the 5th Symposium on Spatial User Interaction*. SUI '17. Brighton, United Kingdom: Association for Computing Machinery, 2017, pp. 99–108. ISBN: 9781450354868. DOI: 10.1145/3131277.3132180. URL: https://doi.org/10.1145/3131277.3132180.

[307] Duc-Minh Pham and Wolfgang Stuerzlinger. "Is the Pen Mightier than the Controller? A Comparison of Input Devices for Selection in Virtual and Augmented Reality". *Proceedings of the 25th ACM Symposium on Virtual Reality Software and Technology*. VRST '19. Parramatta, NSW, Australia: Association for Computing Machinery, 2019. ISBN: 9781450370011. DOI: 10.1145/3359996.3364264. URL: https://doi.org/10.1145/3359996.3364264.

[308] Thammathip Piumsomboon, Adrian Clark, Mark Billinghurst, and Andy Cockburn. "User-Defined Gestures for Augmented Reality". *Conference on Human Factors in Computing Systems - Proceedings*. 2013. ISBN: 9781450318990. DOI: 10.1007/978-3-642-40480-1_18.

[309] Henning Pohl and Kasper Hornbæk. "ElectricItch: Skin Irritation as a Feedback Modality". *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST '18. Berlin, Germany: Association for Computing Machinery, 2018, pp. 765–778. ISBN: 9781450359481. DOI: 10.1145/3242587.3242647. URL: https://doi.org/10.1145/3242587.3242647.

[310] Dominic Potts, Martynas Dabravalskis, and Steven Houben. "TangibleTouch: A Toolkit for Designing Surface-Based Gestures for Tangible Interfaces". *Sixteenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '22. Daejeon, Republic of Korea: Association for Computing Machinery, 2022. ISBN: 9781450391474. DOI: 10.1145/3490149.3502263. URL: https://doi.org/10.1145/3490149.3502263.

[311] Dominic Potts, Kate Loveys, HyunYoung Ha, Shaoyan Huang, Mark Billinghurst, and Elizabeth Broadbent. "ZenG: AR Neurofeedback for Meditative Mixed Reality". *Proceedings of the 2019 on Creativity and Cognition*. C&C '19. San Diego, CA, USA: Association for Computing Machinery, 2019, pp. 583–590. ISBN: 9781450359177. DOI: 10.1145/3325480.3326584. URL: https://doi.org/10.1145/3325480.3326584.

[312] Ivan Poupyrev, Mark Billinghurst, Suzanne Weghorst, and Tadao Ichikawa. "The go-go interaction technique: non-linear mapping for direct manipulation in VR". *Proceedings of the 9th annual ACM symposium on User interface software and technology*. 1996, pp. 79–80.

[313] IVAN POUPYREV and TADAO ICHIKAWA. "Manipulating Objects in Virtual Worlds: Categorization and Empirical Evaluation of Interaction Techniques". *Journal of Visual Languages & Computing* 10.1 (1999), pp. 19–35. ISSN: 1045-926X. DOI: https://doi.org/10.1006/jvlc.1998.0112. URL: https://www.sciencedirect.com/science/article/pii/S1045926X98901124.

[314] Ivan Poupyrev, Tatsushi Nashida, and Makoto Okabe. "Actuation and tangible user interfaces: the Vaucanson duck, robots, and shape displays". *Proceedings of the 1st international conference on Tangible and embedded interaction*. 2007, pp. 205–212.

[315] Ivan Poupyrev, Desney Tan, Mark Billinghurst, Hirokazu Kato, Holger Regenbrecht, and Nobuji Tetsutani. "Tiles: A Mixed Reality Authoring Interface". *(2001) INTERACT 2001 Conference on Human Computer Interaction*. Oct. 2001. URL: https://www.microsoft.com/en-us/research/publication/tiles-mixed-reality-authoring-interface/.

[316] Ivan Poupyrev, Suzanne Weghorst, and Sidney Fels. "Non-lsomorphic 3D rotational techniques". *Conference on Human Factors in Computing Systems - Proceedings*. 2000. ISBN: 1581132166. DOI: 10.1145/332040.332497.

[317] Ivan Poupyrev, Suzanne Weghorst, Takahiro Otsuka, and Tadao Ichikawa. "Amplifying Spatial Rotations in 3D Interfaces". *CHI '99 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '99. Pittsburgh, Pennsylvania: Association for Computing Machinery, 1999, pp. 256–257. ISBN: 1581131585. DOI: 10.1145/632716.632874. URL: https://doi.org/10.1145/632716.632874.

[318] Sara Price, Taciana Pontual Falcão, Jennifer G Sheridan, and George Roussos. "The effect of representation location on interaction in a tangible learning environment". *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*. 2009, pp. 85–92.

[319] Isabel PS Qamar, Rainer Groh, David Holman, and Anne Roudaut. "HCI meets material science: A literature review of morphing materials for the design of shape-changing interfaces". *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 2018, pp. 1–23.

[320] Majken K Rasmussen, Esben W Pedersen, Marianne G Petersen, and Kasper Hornbæk. "Shape-changing interfaces: a review of the design space and open research questions". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2012, pp. 735–744.

[321] Majken K. Rasmussen, Esben W. Pedersen, Marianne G. Petersen, and Kasper Hornbæk. "Shape-Changing Interfaces: A Review of the Design Space and Open Research Questions". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: Association for Computing Machinery, 2012, pp. 735–744. ISBN: 9781450310154. DOI: 10.1145/2207676.2207781. URL: https://doi.org/10.1145/2207676.2207781.

[322] Patrick Reipschläger and Raimund Dachselt. "DesignAR: Immersive 3D-Modeling Combining Augmented Reality with Interactive Displays". *Proceedings of the 2019 ACM International Conference on Interactive Surfaces and Spaces*. ISS '19. Daejeon, Republic of Korea: Association for Computing Machinery, 2019, pp. 29–41. ISBN: 9781450368919. DOI: 10.1145/3343055.3359718. URL: https://doi.org/10.1145/3343055.3359718.

[323] Jun Rekimoto. "Transvision: A hand-held augmented reality system for collaborative design". *Proceeding of Virtual Systems and Multimedia*. Vol. 96. 1996, pp. 18–20.

[324] Julie Rico and Stephen Brewster. "Usable gestures for mobile interfaces: evaluating social acceptability". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2010, pp. 887–896.

[325] Somaiieh Rokhsaritalemi, Abolghasem Sadeghi-Niaraki, and Soo-Mi Choi. "A Review on Mixed Reality: Current Trends, Challenges and Prospects". *Applied Sciences* 10.2 (2020). ISSN: 2076-3417. DOI: 10.3390/app10020636. URL: https://www.mdpi.com/2076-3417/10/2/636.

[326] Claudia Daudén Roquet, Corina Sas, and Dominic Potts. "Exploring Anima: a brain–computer interface for peripheral materialization of mindfulness states during mandala coloring". *Human–Computer Interaction* 0.0 (2021), pp. 1–41. DOI: 10.1080/07370024.2021.1968864. eprint: https://doi.org/10.1080/07370024.2021.1968864. URL: https://doi.org/10.1080/07370024.2021.1968864.

[327] Anne Roudaut, Abhijit Karnik, Markus Löchtefeld, and Sriram Subramanian. “Morphees: toward high” shape resolution” in self-actuated flexible mobile devices”. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2013, pp. 593–602.

[328] *Rubik's Cube*. URL: https://rubiks.com/en-US/productdetails/778988419533. (accessed 15/01/2023).

[329] Robert L Sainburg. “Evidence for a dynamic-dominance hypothesis of handedness”. *Experimental brain research* 142.2 (2002), pp. 241–258.

[330] Kim Sauvé, Dominic Potts, Jason Alexander, and Steven Houben. “A Change of Perspective: How User Orientation Influences the Perception of Physicalizations”. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20. Honolulu, HI, USA: Association for Computing Machinery, 2020, pp. 1–12. ISBN: 9781450367080. DOI: 10.1145/3313831.3376312. URL: https://doi.org/10.1145/3313831.3376312.

[331] Valkyrie Savage, Colin Chang, and Björn Hartmann. “Sauron: Embedded Single-Camera Sensing of Printed Physical User Interfaces”. *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*. UIST '13. St. Andrews, Scotland, United Kingdom: Association for Computing Machinery, 2013, pp. 447–456. ISBN: 9781450322683. DOI: 10.1145/2501988.2501992. URL: https://doi.org/10.1145/2501988.2501992.

[332] Valkyrie Savage, Xiaohan Zhang, and Björn Hartmann. “Midas: Fabricating Custom Capacitive Touch Sensors to Prototype Interactive Objects”. *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology*. UIST '12. Cambridge, Massachusetts, USA: Association for Computing Machinery, 2012, pp. 579–588. ISBN: 9781450315807. DOI: 10.1145/2380116.2380189. URL: https://doi.org/10.1145/2380116.2380189.

[333] Katrin E. Schein and Philipp A. Rauschnabel. “Augmented Reality in Manufacturing: Exploring Workers 2019; Perceptions of Barriers”. *IEEE Transactions on Engineering Management* (2021), pp. 1–14. DOI: 10.1109/TEM.2021.3093833.

[334] Jonas Schjerlund, Kasper Hornbæk, and Joanna Bergström. “Ninja Hands: Using Many Hands to Improve Target Selection in VR”. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. Yokohama, Japan: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445759. URL: https://doi.org/10.1145/3411764.3445759.

[335] Martin Schmitz, Mohammadreza Khalilbeigi, Matthias Balwierz, Roman Lissermann, Max Mühlhäuser, and Jürgen Steimle. "Capricate: A Fabrication Pipeline to Design and 3D Print Capacitive Touch Sensors for Interactive Objects". *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*. UIST '15. Charlotte, NC, USA: Association for Computing Machinery, 2015, pp. 253–258. ISBN: 9781450337793. DOI: 10.1145/2807442.2807503. URL: https://doi.org/10.1145/2807442.2807503.

[336] Stacey D. Scott, M. Sheelagh T. Carpendale, and Kori Inkpen. "Territoriality in Collaborative Tabletop Workspaces". *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work*. CSCW '04. Chicago, Illinois, USA: Association for Computing Machinery, 2004, pp. 294–303. ISBN: 1581138105. DOI: 10.1145/1031607.1031655. URL: https://doi.org/10.1145/1031607.1031655.

[337] Mickael Sereno, Xiyao Wang, Lonni Besançon, Michael J. McGuffin, and Tobias Isenberg. "Collaborative Work in Augmented Reality: A Survey". *IEEE Transactions on Visualization and Computer Graphics* 28.6 (2022), pp. 2530–2549. DOI: 10.1109/TVCG.2020.3032761.

[338] Orit Shaer and Eva Hornecker. *Tangible user interfaces: past, present, and future directions*. Now Publishers Inc, 2010.

[339] Jennifer G Sheridan, Ben W Short, Kristof Van Laerhoven, Nicolas Villar, and Gerd Kortuem. "Exploring cube affordance: Towards a classification of non-verbal dynamics of physical interfaces for wearable computing" (2003).

[340] Jotaro Shigeyama, Takeru Hashimoto, Shigeo Yoshida, Takuji Narumi, Tomohiro Tanikawa, and Michitaka Hirose. "Transcalibur: A Weight Shifting Virtual Reality Controller for 2D Shape Rendering Based on Computational Perception Model". *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. Glasgow, Scotland Uk: Association for Computing Machinery, 2019, pp. 1–11. ISBN: 9781450359702. DOI: 10.1145/3290605.3300241. URL: https://doi.org/10.1145/3290605.3300241.

[341] Ben Shneiderman. "Direct Manipulation: A Step beyond Programming Languages (Abstract Only)". *SIGSOC Bull.* 13.2–3 (May 1981), p. 143. ISSN: 0163-5794. DOI: 10.1145/1015579.810991. URL: https://doi.org/10.1145/1015579.810991.

[342] Ben Shneiderman, Catherine Plaisant, Maxine Cohen, Steven Jacobs, Niklas Elmqvist, and Nicholas Diakopoulos. *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. 6th. Pearson, 2016. ISBN: 013438038X.

[343] Robert W Shumaker, Kristina R Walkup, and Benjamin B Beck. *Animal tool behavior: the use and manufacture of tools by animals*. JHU Press, 2011.

[344] Ludwig Sidenmark, Christopher Clarke, Xuesong Zhang, Jenny Phu, and Hans Gellersen. "Outline Pursuits: Gaze-Assisted Selection of Occluded Objects in Virtual Reality". *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20. Honolulu, HI, USA: Association for Computing Machinery, 2020, pp. 1–13. ISBN: 9781450367080. DOI: 10.1145/3313831.3376438. URL: https://doi.org/10.1145/3313831.3376438.

[345] Ludwig Sidenmark, Dominic Potts, Bill Bapisch, and Hans Gellersen. "Radi-Eye: Hands-Free Radial Interfaces for 3D Interaction Using Gaze-Activated Head-Crossing". *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. Yokohama, Japan: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445697. URL: https://doi.org/10.1145/3411764.3445697.

[346] Gurashish Singh, Alexander Nelson, Ryan Robucci, Chintan Patel, and Nilanjan Banerjee. "Inviz: Low-power personalized gesture recognition using wearable textile capacitive sensor arrays". *2015 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. 2015, pp. 198–206. DOI: 10.1109/PERCOM.2015.7146529.

[347] Mel Slater and Sylvia Wilbur. "A Framework for Immersive Virtual Environments (FIVE): Speculations on the Role of Presence in Virtual Environments". *Presence: Teleoperators and Virtual Environments* 6.6 (Dec. 1997), pp. 603–616. DOI: 10.1162/pres.1997.6.6.603. eprint: https://direct.mit.edu/pvar/article-pdf/6/6/603/1623151/pres.1997.6.6.603.pdf. URL: https://doi.org/10.1162/pres.1997.6.6.603.

[348] *Slinky*. URL: https://www.amazon.co.uk/Slinky-60100-F9L60100-Original/dp/B00000IZKX/ref=sr_1_5?keywords=slinky+toy&qid=1676296603&sr=8-5. (accessed 15/01/2023).

[349] David Canfield Smith, Charles Irby, Ralph Kimball, and Eric Harslem. "The Star User Interface: An Overview". *Proceedings of the June 7-10, 1982, National Computer Conference*. AFIPS '82. Houston, Texas: Association for Computing Machinery, 1982, pp. 515–528. ISBN: 088283035X. DOI: 10.1145/1500774.1500840. URL: https://doi.org/10.1145/1500774.1500840.

[350] Graham Smith, Tim Salzman, and Wolfgang Stuerzlinger. "3D Scene Manipulation with 2D Devices and Constraints". *Proceedings of Graphics Interface 2001*. GI '01. Ottawa, Ontario, Canada: Canadian Information Processing Society, 2001, pp. 135–142. ISBN: 0968880800.

[351] Graham Smith and Wolfgang Stuerzlinger. "Integration of Constraints into a VR Environment". *VRIC, Virtual Reality International Conference*. 2001.

[352] Rajinder Sodhi, Ivan Poupyrev, Matthew Glisson, and Ali Israr. "AIREAL: Interactive Tactile Experiences in Free Air". *ACM Trans. Graph.* 32.4 (July 2013). ISSN: 0730-0301. DOI: 10.1145/2461912.2462007. URL: https://doi.org/10.1145/2461912.2462007.

[353] Peng Song, Wooi Boon Goh, William Hutama, Chi Wing Fu, and Xiaopei Liu. "A handle bar metaphor for virtual object manipulation with mid-air interaction". *Conference on Human Factors in Computing Systems - Proceedings*. 2012. ISBN: 9781450310154. DOI: 10.1145/2207676.2208585.

[354] Maximilian Speicher, Brian D. Hall, and Michael Nebeling. "What is Mixed Reality?" *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. Glasgow, Scotland Uk: Association for Computing Machinery, 2019, pp. 1–15. ISBN: 9781450359702. DOI: 10.1145/3290605.3300767. URL: https://doi.org/10.1145/3290605.3300767.

[355] Katta Spiel. "The Bodies of TEI–Investigating Norms and Assumptions in the Design of Embodied Interaction". *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction*. 2021, pp. 1–19.

[356] *Stanford Bunny*. URL: http://graphics.stanford.edu/data/3Dscanrep/. (accessed 24/01/2023).

[357] Christian Stetco, Stephan Mühlbacher-Karrer, Matteo Lucchi, Matthias Weyrer, Lisa-Marie Faller, and Hubert Zangl. "Gesture-based Contactless Control of Mobile Manipulators using Capacitive Sensing". *2020 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*. 2020, pp. 1–6. DOI: 10.1109/I2MTC43012.2020.9128751.

[358] Richard Stoakley, Matthew J Conway, and Randy Pausch. "Virtual reality on a WIM: interactive worlds in miniature". *Proceedings of the SIGCHI conference on Human factors in computing systems*. 1995, pp. 265–272.

[359] Patrick L. Strandholt, Oana A. Dogaru, Niels C. Nilsson, Rolf Nordahl, and Stefania Serafin. "Knock on Wood: Combining Redirected Touching and Physical Props for Tool-Based Interaction in Virtual Reality". *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20. Honolulu, HI, USA: Association for Computing Machinery, 2020, pp. 1–13. ISBN: 9781450367080. DOI: 10.1145/3313831.3376303. URL: https://doi.org/10.1145/3313831.3376303.

[360] *Stretchy String*. URL: https://www.amazon.co.uk/KeNeer-Sensory-Fidget-Stretchy-Anxiety/dp/B075CFKQY6. (accessed 15/01/2023).

[361] Miriam Sturdee and Jason Alexander. "Analysis and Classification of Shape-Changing Interfaces for Design and Application-Based Research". *ACM Comput. Surv.* 51.1 (Jan. 2018). ISSN: 0360-0300. DOI: 10.1145/3143559. URL: https://doi.org/10.1145/3143559.

[362] Miriam Sturdee, John Hardy, Nick Dunn, and Jason Alexander. "A Public Ideation of Shape-Changing Applications". *Proceedings of the 2015 International Conference on Interactive Tabletops &; Surfaces*. ITS '15. Madeira, Portugal: Association for Computing Machinery, 2015, pp. 219–228. ISBN: 9781450338998. DOI: 10.1145/2817721.2817734. URL: https://doi.org/10.1145/2817721.2817734.

[363] D.J. Sturman and D. Zeltzer. "A survey of glove-based input". *IEEE Computer Graphics and Applications* 14.1 (1994), pp. 30–39. DOI: 10.1109/38.250916.

[364] Kenji Sugihara, Mai Otsuki, Asako Kimura, Fumihisa Shibata, and Hideyuki Tamura. "MAI Painting Brush++: Augmenting the Feeling of Painting with New Visual and Tactile Feedback Mechanisms". *Proceedings of the 24th Annual ACM Symposium Adjunct on User Interface Software and Technology*. UIST '11 Adjunct. Santa Barbara, California, USA: Association for Computing Machinery, 2011, pp. 13–14. ISBN: 9781450310147. DOI: 10.1145/2046396.2046404. URL: https://doi.org/10.1145/2046396.2046404.

[365] Junwei Sun, Wolfgang Stuerzlinger, and Bernhard E. Riecke. "Comparing Input Methods and Cursors for 3D Positioning with Head-Mounted Displays". *Proceedings of the 15th ACM Symposium on Applied Perception*. SAP '18. Vancouver, British Columbia, Canada: Association for Computing Machinery, 2018. ISBN: 9781450358941. DOI: 10.1145/3225153.3225167. URL: https://doi.org/10.1145/3225153.3225167.

[366] Junwei Sun, Wolfgang Stuerzlinger, and Dmitri Shuralyov. "SHIFT-Sliding and DEPTH-POP for 3D Positioning". *Proceedings of the 2016 Symposium on Spatial User Interaction*. SUI '16. Tokyo, Japan: Association for Computing Machinery, 2016, pp. 69–78. ISBN: 9781450340687. DOI: 10.1145/2983310.2985748. URL: https://doi.org/10.1145/2983310.2985748.

[367] Yuqian Sun, Shigeo Yoshida, Takuji Narumi, and Michitaka Hirose. "Handheld Haptic Interface for Rendering Size, Shape, and Stiffness of Virtual Objects". *Proceedings of the 2018 ACM International Conference on Interactive Surfaces and Spaces*. ISS '18. Tokyo, Japan: Association for Computing Machinery, 2018, pp. 411–414. ISBN: 9781450356947. DOI: 10.1145/3279778.3281759. URL: https://doi.org/10.1145/3279778.3281759.

[368] Ivan Sutherland. "The ultimate display" (1965).

[369] Ivan E. Sutherland. "A Head-Mounted Three Dimensional Display". *Proceedings of the December 9-11, 1968, Fall Joint Computer Conference, Part I*. AFIPS '68 (Fall, part I). San Francisco, California: Association for Computing Machinery, 1968, pp. 757–764. ISBN: 9781450378994. DOI: 10.1145/1476589.1476686. URL: https://doi.org/10.1145/1476589.1476686.

[370] J. Edward Swan, Gurjot Singh, and Stephen R. Ellis. "Matching and Reaching Depth Judgments with Real and Augmented Reality Targets". *IEEE Transactions on Visualization and Computer Graphics* 21.11 (2015), pp. 1289–1298. DOI: 10.1109/TVCG.2015.2459895.

[371] Zsolt Szalavári, Dieter Schmalstieg, Anton Fuhrmann, and Michael Gervautz. ""Studierstube": An environment for collaboration in augmented reality". *Virtual Reality* 3.1 (1998), pp. 37–48.

[372] Faisal Taher, John Hardy, Abhijit Karnik, Christian Weichel, Yvonne Jansen, Kasper Hornbæk, and Jason Alexander. "Exploring Interactions with Physically Dynamic Bar Charts". *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. CHI '15. Seoul, Republic of Korea: Association for Computing Machinery, 2015, pp. 3237–3246. ISBN: 9781450331456. DOI: 10.1145/2702123.2702604. URL: https://doi.org/10.1145/2702123.2702604.

[373] Faisal Taher, John Vidler, and Jason Alexander. "A Characterization of Actuation Techniques for Generating Movement in Shape-Changing Interfaces". *International Journal of Human–Computer Interaction* 33.5 (2017), pp. 385–398. DOI: 10.1080/10447318.2016.1250372. eprint: https://doi.org/10.1080/10447318.2016.1250372. URL: https://doi.org/10.1080/10447318.2016.1250372.

[374] Takafumi Taketomi, Hideaki Uchiyama, and Sei Ikeda. "Visual SLAM algorithms: A survey from 2010 to 2016". *IPSJ Transactions on Computer Vision and Applications* 9.1 (2017), pp. 1–11.

[375] D. S. Tan, G. G. Robertson, and M. Czerwinski. "Exploring 3D navigation: Combining speed-coupled flying with orbiting". *Conference on Human Factors in Computing Systems - Proceedings*. 2001.

[376] OpenCV team. *OpenCV.org*. https://opencv.org/. 2023.

[377] Robert J. Teather, Andriy Pavlovych, Wolfgang Stuerzlinger, and I. Scott MacKenzie. "Effects of tracking technology, latency, and spatial jitter on object movement". *2009 IEEE Symposium on 3D User Interfaces*. 2009, pp. 43–50. DOI: 10.1109/3DUI.2009.4811204.

[378] Shan-Yuan Teng, Tzu-Sheng Kuo, Chi Wang, Chi-huan Chiang, Da-Yuan Huang, Liwei Chan, and Bing-Yu Chen. “PuPoP: Pop-up Prop on Palm for Virtual Reality”. *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST ’18. Berlin, Germany: Association for Computing Machinery, 2018, pp. 5–17. ISBN: 9781450359481. DOI: 10.1145/3242587.3242628. URL: https://doi.org/10.1145/3242587.3242628.

[379] *Think Link*. URL: https://www.etsy.com/uk/listing/491319961/think-link-fidget-toy-multiple-colors. (accessed 15/01/2023).

[380] *Thinking Putty*. URL: https://www.amazon.co.uk/Magnet-Expert-Magnetic-Putty-Cube/dp/B006FTCL02/ref=sr_1_47?crid=NXZYV29UKARG&keywords=putty&qid=1676298466&sprefix=putty%5C%2Caps%5C%2C77&sr=8-47. (accessed 15/01/2023).

[381] Edward Thompson, Dominic Potts, John Hardy, Barry Porter, and Steven Houben. “AmbiDots: An Ambient Interface to Mediate Casual Social Settings through Peripheral Interaction”. *Proceedings of the 33rd Australian Conference on Human-Computer Interaction*. OzCHI ’21. Melbourne, VIC, Australia: Association for Computing Machinery, 2022, pp. 99–110. ISBN: 9781450395984. DOI: 10.1145/3520495.3520504. URL: https://doi.org/10.1145/3520495.3520504.

[382] *Tilt Brush by Google*. URL: https://www.tiltbrush.com/. (accessed 20/01/2023).

[383] Eleftherios Triantafyllidis and Zhibin Li. “The Challenges in Modeling Human Performance in 3D Space with Fitts’ Law”. CHI EA ’21. Yokohama, Japan: Association for Computing Machinery, 2021. ISBN: 9781450380959. DOI: 10.1145/3411763.3443442. URL: https://doi.org/10.1145/3411763.3443442.

[384] Giovanni Maria Troiano, Esben Warming Pedersen, and Kasper Hornbæk. “User-defined gestures for elastic, deformable displays”. *Proceedings of the Workshop on Advanced Visual Interfaces AVI*. 2014. ISBN: 9781450327756. DOI: 10.1145/2598153.2598184.

[385] Theophanis Tsandilas. “Fallacies of Agreement: A Critical Review of Consensus Assessment Methods for Gesture Elicitation”. *ACM Trans. Comput.-Hum. Interact.* 25.3 (June 2018). ISSN: 1073-0516. DOI: 10.1145/3182168. URL: https://doi.org/10.1145/3182168.

[386] Hideaki Uchiyama and Eric Marchand. “Object Detection and Pose Tracking for Augmented Reality: Recent Approaches”. *18th Korea-Japan Joint Workshop on Frontiers of Computer Vision (FCV)*. Kawasaki, Japan, Feb. 2012. URL: https://hal.inria.fr/hal-00751704.

[387] Brygg Ullmer and Hiroshi Ishii. "Emerging frameworks for tangible user interfaces". *IBM systems journal* 39.3.4 (2000), pp. 915–931.

[388] Brygg Ullmer, Hiroshi Ishii, and Dylan Glas. "mediaBlocks: physical containers, transports, and controls for online media". *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. 1998, pp. 379–386.

[389] Brygg Ullmer, Hiroshi Ishii, and Robert J. K. Jacob. "Token+constraint Systems for Tangible Interaction with Digital Information". *ACM Trans. Comput.-Hum. Interact.* 12.1 (Mar. 2005), pp. 81–118. ISSN: 1073-0516. DOI: 10.1145/1057237.1057242. URL: https://doi.org/10.1145/1057237.1057242.

[390] *Ultra Leap Develop Mount for Integrated Hand Tracking in VR/AR*. URL: https://www.ultraleap.com/product/vr-developer-mount/. (accessed 20/01/2023).

[391] *Valve: Index*. URL: https://store.steampowered.com/sub/354231/. (accessed 15/01/2023).

[392] Jelle Van Dijk, Remko Van Der Lugt, and Caroline Hummels. "Beyond distributed representation: embodied cognition design supporting socio-sensorimotor couplings". *Proceedings of the 8th International Conference on Tangible, Embedded and Embodied Interaction*. 2014, pp. 181–188.

[393] Kristof Van Laerhoven, Nicolas Villar, Albrecht Schmidt, Gerd Kortuem, and Hans Gellersen. "Using an Autonomous Cube for Basic Navigation and Input". *Proceedings of the 5th International Conference on Multimodal Interfaces*. ICMI '03. Vancouver, British Columbia, Canada: Association for Computing Machinery, 2003, pp. 203–210. ISBN: 1581136218. DOI: 10.1145/958432.958472. URL: https://doi.org/10.1145/958432.958472.

[394] Lode Vanacken, Tovi Grossman, and Karin Coninx. "Exploring the Effects of Environment Density and Target Visibility on Object Selection in 3D Virtual Environments". *2007 IEEE Symposium on 3D User Interfaces*. 2007. DOI: 10.1109/3DUI.2007.340783.

[395] Lode Vanacken, Tovi Grossman, and Karin Coninx. "Multimodal selection techniques for dense and occluded 3D virtual environments". *International Journal of Human-Computer Studies* 67.3 (2009). Current trends in 3D user interface research, pp. 237–255. ISSN: 1071-5819. DOI: https://doi.org/10.1016/j.ijhcs.2008.09.001. URL: https://www.sciencedirect.com/science/article/pii/S1071581908001109.

[396] Virag Varga, Gergely Vakulya, Benjamin Buergisser, Nathan Riopelle, Fabio Zund, Robert W. Sumner, Thomas R. Gross, and Alanson Sample. "Real-Time Capture of Holistic Tangible Interactions". *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '21. Salzburg, Austria: Association for Computing Machinery, 2021. ISBN: 9781450382137. DOI: 10.1145/3430524.3440658. URL: https://doi.org/10.1145/3430524.3440658.

[397] Radu-Daniel Vatavu and Jacob O Wobbrock. "Formalizing agreement analysis for elicitation studies: new measures, significance test, and toolkit". *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 2015, pp. 1325–1334.

[398] Cati Vaucelle and Hiroshi Ishii. "Picture This! Film assembly using toy gestures". *Proceedings of the 10th international conference on Ubiquitous computing*. 2008, pp. 350–359.

[399] Aaron Visschedijk, Hyunyoung Kim, Carlos Tejada, and Daniel Ashbrook. "ClipWidgets: 3D-Printed Modular Tangible UI Extensions for Smartphones". *Sixteenth International Conference on Tangible, Embedded, and Embodied Interaction*. TEI '22. Daejeon, Republic of Korea: Association for Computing Machinery, 2022. ISBN: 9781450391474. DOI: 10.1145/3490149.3501314. URL: https://doi.org/10.1145/3490149.3501314.

[400] Panagiotis Vogiatzidakis and Panayiotis Koutsabasis. "Gesture Elicitation Studies for Mid-Air Interaction: A Review". *Multimodal Technologies and Interaction* (2018). ISSN: 2414-4088. DOI: 10.3390/mti2040065.

[401] Tijana Vuletic, Alex Duffy, Laura Hay, Chris Patrick McTeague, Gerard Campbell, Pei Ling Choo, Madeleine Grealy, et al. "Natural and intuitive gesture interaction for 3D object manipulation in conceptual design". *Proceedings of International Design Conference, DESIGN*. 2018. ISBN: 9789537738594. DOI: 10.21278/idc.2018.0321.

[402] *Wacky Tracks*. URL: https://www.additudemag.com/product/neliblu-wacky-tracks-snap-and-click/. (accessed 15/01/2023).

[403] James A Walsh, Stewart von Itzstein, and Bruce H Thomas. "Tangible Agile Mapping: Ad-hoc Tangible User Interaction Definition". *Proceedings of the Fourteenth Australasian User Interface Conference - Volume 139*. 2013. ISBN: 978-1-921770-24-1.

[404] Zhimin Wang, Haofei Wang, Huangyue Yu, and Feng Lu. "Interaction With Gaze, Gesture, and Speech in a Flexibly Configurable Augmented Reality System". *IEEE Transactions on Human-Machine Systems* 51.5 (2021), pp. 524–534. DOI: 10.1109/THMS.2021.3097973.

[405] Y. Watanabe, A. Cassinelli, T. Komuro, and M. Ishikawa. "The deformable workspace: A membrane between real and virtual space". *2008 3rd IEEE International Workshop on Horizontal Interactive Human Computer Systems*. 2008, pp. 145–152.

[406] Christian Weichel, Manfred Lau, David Kim, Nicolas Villar, and Hans W. Gellersen. "MixFab: A Mixed-Reality Environment for Personal Fabrication". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '14. Toronto, Ontario, Canada: Association for Computing Machinery, 2014, pp. 3855–3864. ISBN: 9781450324731. DOI: 10.1145/2556288.2557090. URL: https://doi.org/10.1145/2556288.2557090.

[407] Malte Weiss, Florian Schwarz, Simon Jakubowski, and Jan Borchers. "Madgets: Actuating Widgets on Interactive Tabletops". *Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology*. UIST '10. New York, New York, USA: Association for Computing Machinery, 2010, pp. 293–302. ISBN: 9781450302715. DOI: 10.1145/1866029.1866075. URL: https://doi.org/10.1145/1866029.1866075.

[408] Malte Weiss, Julie Wagner, Yvonne Jansen, Roger Jennings, Ramsin Khoshabeh, James D. Hollan, and Jan Borchers. "SLAP Widgets: Bridging the Gap between Virtual and Physical Controls on Tabletops". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '09. Boston, MA, USA: Association for Computing Machinery, 2009, pp. 481–490. ISBN: 9781605582467. DOI: 10.1145/1518701.1518779. URL: https://doi.org/10.1145/1518701.1518779.

[409] Greg Welch, Gary Bishop, Leandra Vicci, Stephen Brumback, Kurtis Keller, and D'nardo Colucci. "The HiBall Tracker: High-Performance Wide-Area Tracking for Virtual and Augmented Environments". *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*. VRST '99. London, United Kingdom: Association for Computing Machinery, 1999, 1–ff. ISBN: 1581131410. DOI: 10.1145/323663.323664. URL: https://doi.org/10.1145/323663.323664.

[410] Thomas Wells, Dominic Potts, and Steven Houben. "A Study into the Effect of Mobile Device Configurations on Co-Located Collaboration Using AR". *Proc. ACM Hum.-Comput. Interact.* 6.MHCI (Sept. 2022). DOI: 10.1145/3546735. URL: https://doi.org/10.1145/3546735.

[411] Gerold Wesche and Hans-Peter Seidel. "FreeDrawer: A Free-Form Sketching System on the Responsive Workbench". *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*. VRST '01. Baniff, Alberta, Canada: Association for Computing Machinery, 2001, pp. 167–174. ISBN: 1581134274.

DOI: 10.1145/505008.505041. URL: https://doi.org/10.1145/505008.505041.

[412] Eric Whitmire, Hrvoje Benko, Christian Holz, Eyal Ofek, and Mike Sinclair. "Haptic Revolver: Touch, Shear, Texture, and Shape Rendering on a Reconfigurable Virtual Reality Controller". *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: Association for Computing Machinery, 2018, pp. 1–12. ISBN: 9781450356206. DOI: 10.1145/3173574.3173660. URL: https://doi.org/10.1145/3173574.3173660.

[413] Daniel Wigdor, Hrvoje Benko, John Pella, Jarrod Lombardo, and Sarah Williams. "Rock & Rails: Extending Multi-Touch Interactions with Shape Gestures to Enable Precise Spatial Manipulations". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '11. Vancouver, BC, Canada: Association for Computing Machinery, 2011, pp. 1581–1590. ISBN: 9781450302289. DOI: 10.1145/1978942.1979173. URL: https://doi.org/10.1145/1978942.1979173.

[414] Adam S. Williams, Jason Garcia, and Francisco Ortega. "Understanding Multimodal User Gesture and Speech Behavior for Object Manipulation in Augmented Reality Using Elicitation". *IEEE Transactions on Visualization and Computer Graphics* 26.12 (2020), pp. 3479–3489. DOI: 10.1109/TVCG.2020.3023566.

[415] Frank R Wilson. *The hand: How its use shapes the brain, language, and human culture*. Vintage, 1999.

[416] Hanno Wirtz, Jan Rüth, Martin Serror, Jó Ágila Bitsch Link, and Klaus Wehrle. "Opportunistic Interaction in the Challenged Internet of Things". *Proceedings of the 9th ACM MobiCom Workshop on Challenged Networks*. CHANTS '14. Maui, Hawaii, USA: Association for Computing Machinery, 2014, pp. 7–12. ISBN: 9781450330718. DOI: 10.1145/2645672.2645679. URL: https://doi.org/10.1145/2645672.2645679.

[417] Anusha Withana, Daniel Groeger, and Jürgen Steimle. "Tacttoo: A Thin and Feel-Through Tattoo for On-Skin Tactile Output". *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST '18. Berlin, Germany: Association for Computing Machinery, 2018, pp. 365–378. ISBN: 9781450359481. DOI: 10.1145/3242587.3242645. URL: https://doi.org/10.1145/3242587.3242645.

[418] Jacob O. Wobbrock, Htet Htet Aung, Brandon Rothrock, and Brad A. Myers. "Maximizing the guessability of symbolic input". *Conference on Human Factors in Computing Systems - Proceedings*. 2005. ISBN: 1595930027. DOI: 10.1145/1056808.1057043.

[419] Jacob O. Wobbrock, Meredith Ringel Morris, and Andrew D. Wilson. "User-defined gestures for surface computing". *Conference on Human Factors in Computing Systems - Proceedings*. 2009. ISBN: 9781605582474. DOI: 10.1145/1518701.1518866.

[420] Hsin-Kai Wu, Silvia Wen-Yu Lee, Hsin-Yi Chang, and Jyh-Chong Liang. "Current status, opportunities and challenges of augmented reality in education". *Computers & Education* 62 (2013), pp. 41–49. ISSN: 0360-1315. DOI: https://doi.org/10.1016/j.compedu.2012.10.024. URL: https://www.sciencedirect.com/science/article/pii/S0360131512002527.

[421] Yiqi Xiao and Renke He. "The intuitive grasp interface: design and evaluation of micro-gestures on the steering wheel for driving scenario". *Universal Access in the Information Society* (2019), pp. 1–18.

[422] Kotaro Yamaguchi, Ginga Kato, Yoshihiro Kuroda, Kiyoshi Kiyokawa, and Haruo Takemura. "A Non-Grounded and Encountered-Type Haptic Display Using a Drone". *Proceedings of the 2016 Symposium on Spatial User Interaction*. SUI '16. Tokyo, Japan: Association for Computing Machinery, 2016, pp. 43–46. ISBN: 9781450340687. DOI: 10.1145/2983310.2985746. URL: https://doi.org/10.1145/2983310.2985746.

[423] Jackie (Junrui) Yang, Hiroshi Horii, Alexander Thayer, and Rafael Ballagas. "VR Grabbers: Ungrounded Haptic Retargeting for Precision Grabbing Tools". *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. UIST '18. Berlin, Germany: Association for Computing Machinery, 2018, pp. 889–899. ISBN: 9781450359481. DOI: 10.1145/3242587.3242643. URL: https://doi.org/10.1145/3242587.3242643.

[424] Hui-Shyong Yeo, Ryosuke Minami, Kirill Rodriguez, George Shaker, and Aaron Quigley. "Exploring Tangible Interactions with Radar Sensing". *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2.4 (Dec. 2018). DOI: 10.1145/3287078. URL: https://doi.org/10.1145/3287078.

[425] Difeng Yu, Xueshi Lu, Rongkai Shi, Hai-Ning Liang, Tilman Dingler, Eduardo Velloso, and Jorge Goncalves. "Gaze-Supported 3D Object Manipulation in Virtual Reality". *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. Yokohama, Japan: Association for Computing Machinery, 2021. ISBN: 9781450380966. DOI: 10.1145/3411764.3445343. URL: https://doi.org/10.1145/3411764.3445343.

[426] Difeng Yu, Qiushi Zhou, Tilman Dingler, Eduardo Velloso, and Jorge Goncalves. "Blending On-Body and Mid-Air Interaction in Virtual Reality". *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. 2022, pp. 637–646. DOI: 10.1109/ISMAR55827.2022.00081.

[427] Difeng Yu, Qiushi Zhou, Joshua Newn, Tilman Dingler, Eduardo Velloso, and Jorge Goncalves. "Fully-Occluded Target Selection in Virtual Reality". *IEEE Transactions on Visualization and Computer Graphics* 26.12 (2020), pp. 3402–3413. DOI: 10.1109/TVCG.2020.3023606.

[428] Robert C. Zeleznik, Andrew S. Forsberg, and Paul S. Strauss. "Two Pointer Input for 3D Interaction". *Proceedings of the 1997 Symposium on Interactive 3D Graphics*. I3D '97. Providence, Rhode Island, USA: Association for Computing Machinery, 1997, pp. 115–120. ISBN: 0897918843. DOI: 10.1145/253284.253316. URL: https://doi.org/10.1145/253284.253316.

[429] Shumin Zhai. "User Performance in Relation to 3D Input Device Design". *SIGGRAPH Comput. Graph.* 32.4 (Nov. 1998), pp. 50–54. ISSN: 0097-8930. DOI: 10.1145/307710.307728. URL: https://doi.org/10.1145/307710.307728.

[430] Shumin Zhai, William Buxton, and Paul Milgram. "The "silk cursor" investigating transparency for 3d target acquisition". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1994, pp. 459–464.

[431] Clement Zheng, Peter Gyory, and Ellen Yi-Luen Do. "Tangible Interfaces with Printed Paper Markers". *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. DIS '20. Eindhoven, Netherlands: Association for Computing Machinery, 2020, pp. 909–923. ISBN: 9781450369749. DOI: 10.1145/3357236.3395578. URL: https://doi.org/10.1145/3357236.3395578.

[432] Qian Zhou, George Fitzmaurice, and Fraser Anderson. "In-Depth Mouse: Integrating Desktop Mouse into Virtual Reality". *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. CHI '22. New Orleans, LA, USA: Association for Computing Machinery, 2022. ISBN: 9781450391573. DOI: 10.1145/3491102.3501884. URL: https://doi.org/10.1145/3491102.3501884.

[433] Egui Zhu, Arash Hadadgar, Italo Masiello, and Nabil Zary. "Augmented reality in healthcare education: an integrative review". *PeerJ* 2 (2014), e469.

[434] Kening Zhu, Taizhou Chen, Feng Han, and Yi Shiun Wu. "HaptWist: Creating Interactive Haptic Proxies in Virtual Reality Using Low-cost Twistable Artefacts". *Conference on Human Factors in Computing Systems - Proceedings*. 2019. ISBN: 9781450359702. DOI: 10.1145/3290605.3300923.