Quantifying Epistemic Actions in Human-Computer Interaction

DOCTORAL THESIS

Augusto Emanuel Abreu Esteves
DOCTORATE IN INFORMATICS ENGINEERING
SPECIALTY: HUMAN-COMPUTER INTERACTION

UNIVERSIDADE da MADEIRA
A Nossa Universidade
www.uma.pt

January | 2015

FCT Fundação para a Ciência e a Tecnologia
MINISTÉRIO DA EDUCAÇÃO E CIÊNCIA

Quantifying Epistemic Actions in Human-Computer Interaction

DOCTORAL THESIS

Augusto Emanuel Abreu Esteves
DOCTORATE IN INFORMATICS ENGINEERING
SPECIALTY: HUMAN-COMPUTER INTERACTION

SUPERVISOR
Ian Oakley

UNIVERSIDADE DA MADEIRA

PhD Dissertation

UNDERSTANDING EPISTEMIC ACTIONS IN HUMAN-COMPUTER INTERACTION

AUGUSTO ESTEVES

IAN OAKLEY

Ulsan National Institute of Science and Technology

ELISE VAN DEN HOVEN

University of Technology, Sydney

Technische Universiteit Eindhoven

NUNO CORREIA

Universidade Nova de Lisboa

2043904@student.uma.pt • +351 963005724 • mysecondplace.org

To my wife, Hana
And my little brothers
Rodrigo and Simão.

ACKNOWLEDGEMENTS

I dedicate this dissertation to the various individuals that encouraged, inspired and supported me throughout these last four years. I would like to start with my supervisor, Professor Ian Oakley, for believing in me and encouraging me to pursue a PhD – my life would have been tremendously different if not for you.

To Professor Nuno Nunes, for convincing me to pursue my studies at the University of Madeira (UMa) and continuing to play a pivotal role in the development and growth of my research field at our institution. To my brilliant coworkers at UMa – Michelle Scott, Filipe Quintal, Mary Barreto, Lucas Pereira and Jayant Venkatanathan – for your feedback, advice and unforgettable friendship.

To Andrea Bianchi, for having always found the time in his hectic schedule to help and assist me while I was visiting KAIST. To Elise van den Hoven, Saskia Bakker, Jacquelyn van Kampen, Josha Woudstra, Arne Wessels and Nikos Batalas, for welcoming me in your lives at Eindhoven and truly making me feel at home. To all my coworkers at the Color Lab in UNIST, specially Jun Daeun and Kim Seong Hye, for having the kindness and patience to accommodate me and all of my questions about your wonderful culture.

I also dedicate this dissertation to my wonderful family, which has doubled in size since the start of my PhD. They provide me with endless love and affection, and make me a better man. I have overcome all the struggles and difficulties in my life because of you.

Finally, I thank all my co-authors and the participants that took the time to help me with the research I present in this dissertation. This work was supported by the Portuguese Foundation of Science and Technology (FCT), under the SFRH/BD70203/2010 PhD scholarship.

PUBLISHED WORK

The majority of the work presented in this thesis results of a series of publications in several SIGCHI and Association for Computing Machinery (ACM) peer-reviewed conferences, between 2011 and 2015.

Dissertation Chapters

Chapter Four

Esteves, A. and Oakley, I. 2015. tiREC: A Tool For The Recording And Sharing Of Quantitative Data From Tangible Systems. In *submission at 15th International Conference on Human-Computer Interaction* (INTERACT '15). ACM.

Esteves, A., Quintal, F. and Oakley, I. 2013. jamTable: Can Physical Interfaces Support the Collaboration between Novice and Experienced Musicians? In *I. Oakley and S. Brewster (eds.)* Haptic and Audio Interaction Design, 7989, 99-108.

Esteves, A. & Oakley I. 2011. Informing Design by Recording Tangible Interaction. In *Proceedings of the Extended Abstracts of the 29th Annual ACM Conference on Human Factors in Computing Systems (CHI '11)*, ACM, New York, NY, USA, 2077-2082.

Chapter Five

Esteves, A., Hoven, E. van den and Oakley I. 2013. Physical Games or Digital Games? Comparing Support for Mental Projection in Tangible and Virtual Representations of a Problem Solving Task. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction* (TEI '13). ACM, New York, NY, USA, 167-174.

Chapter Six

Esteves, A., Scott, M. and Oakley I. 2013. Designing to Support Offline Activities on Interactive Surfaces. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction* (TEI '13). ACM, New York, NY, USA, 147-154.

Esteves, A. & Oakley I. 2011. Design for interface consistency or embodied facilitation? In *adjunct Proceedings of the 29th Annual ACM Conference on Human Factors in Computing Systems* (CHI '11), Workshop on Embodied Interaction: Theory & Practice in HCI, 37-40.

Esteves, A. & Oakley I. 2011. Eco Planner: A Tabletop System for Scheduling Sustainable Routines. In *Extended Abstracts of the 8th International Conference on Tangible, Embedded and Embodied Interaction* (TEI '11).

Chapter Seven

Esteves, A., Bakker, S., Antle, A., May, A., Warren, J. and Oakley, I. 2015. The ATB Framework: Quantifying and Classifying Epistemic Strategies in Tangible Problem-Solving Tasks. In *Proceedings of the 9th International Conference on Tangible, Embedded and Embodied Interaction* (TEI '15). ACM, New York, NY, USA.

Esteves, A., Bakker, S., Antle, A., May, A., Warren, J. and Oakley, I. 2014. Classifying Physical Strategies in Tangible Tasks: A Video-Coding Framework For Epistemic Actions. In *Proceedings of the Extended Abstracts of the 32nd Annual ACM Conference on Human Factors in Computing Systems* (CHI '14), ACM, New York, USA, 1843-1848.

Other Publications

Esteves, A., Quintal, F. and Oakley, I. 2014. TouchCloud: An Exploratory Study in Physically Tagging the Cloud. In *Extended Abstracts of the 8th International Conference on Tangible, Embedded and Embodied Interaction* (TEI '14).

Oakley, I. & Esteves, A. 2013. On and Offline Tangible Interaction: Studying the Secret Lives of Augmented Objects. In *International Symposium on Ubiquitous Virtual Reality* (ISUVR 2013), 5-6.

Augusto Esteves. 2012. Designing tangible interaction for embodied facilitation. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction* (TEI '12), Stephen N. Spencer (Ed.). ACM, New York, NY, USA, 395-396.

DISSERTATION STATEMENT

Embodied cognition is a rich and appropriate source of valuable quantitative perspectives that can inform the design and evaluation of tangible interaction.

ABSTRACT

As digital systems move away from traditional desktop setups, new interaction paradigms are emerging that better integrate with users' real-world surroundings, and better support users' individual needs. While promising, these modern interaction paradigms also present new challenges, such as a lack of paradigm-specific tools to systematically evaluate and fully understand their use. This dissertation tackles this issue by framing empirical studies of three novel digital systems in embodied cognition – an exciting new perspective in cognitive science where the body and its interactions with the *physical* world take a central role in human cognition. This is achieved by first, focusing the design of all these systems on a contemporary interaction paradigm that emphasizes *physical* interaction on tangible interaction, a contemporary interaction paradigm; and second, by comprehensively studying user performance in these systems through a set of novel performance metrics grounded on **epistemic actions**, a relatively well established and studied construct in the literature on embodied cognition. The first system presented in this dissertation is an augmented Four-in-a-row board game. Three different versions of the game were developed, based on three different interaction paradigms (tangible, touch and mouse), and a repeated measures study involving 36 participants measured the occurrence of three simple epistemic actions across these three interfaces. The results

highlight the relevance of epistemic actions in such a task and suggest that the different interaction paradigms afford instantiation of these actions in different ways. Additionally, the tangible version of the system supports the most rapid execution of these actions, providing novel quantitative insights into the real benefits of tangible systems. The second system presented in this dissertation is a tangible tabletop scheduling application. Two studies with single and paired users provide several insights into the impact of epistemic actions on the user experience when these are performed outside of a system's sensing boundaries. These insights are clustered by the form, size and location of ideal interface areas for such offline epistemic actions to occur, as well as how can physical tokens be designed to better support them. Finally, and based on the results obtained to this point, the last study presented in this dissertation directly addresses the lack of empirical tools to formally evaluate tangible interaction. It presents a video-coding framework grounded on a systematic literature review of 78 papers, and evaluates its value as metric through a 60 participant study performed across three different research laboratories. The results highlight the usefulness and power of epistemic actions as a performance metric for tangible systems. In sum, through the use of such novel metrics in each of the three studies presented, this dissertation provides a better understanding of the real impact and benefits of designing and developing systems that feature tangible interaction.

TABLE OF CONTENTS

FIGURES

TABLES

INTRODUCTION

1 Technological advances over the last 30 years have unchained human-computer interaction (HCI) research from desktop computers rooted on mouse and keyboard input, and enabled the study of digital systems that permeate almost every facet of human experience. This has led to the development of numerous concepts and interaction paradigms within HCI over the years, including multimodal and tangible interaction, brain-computer interfaces, ubiquitous computing (e.g., the internet-of-things, wearables, ambient intelligence) and virtual and augmented reality. While these concepts and paradigms are ripe with promise, such as of "calm technology" [WB96] or of digital systems that are more intuitive [e.g., SH10], responsive [e.g., VWH05] or adaptive [e.g., AW09], they are equally ripe with challenges. One of these challenges is in creating a paradigm-specific body of knowledge of how to design, develop, and evaluate these modern digital systems. This dissertation focuses specifically on evaluation, as it acknowledges that many of the proposed theoretical benefits introduced by these interaction paradigms still require empirical validation [e.g., AM00, O99, ZAM+12]. However this is not easy to achieve. Many of the evaluation techniques used in HCI are tailored for classical graphical user interfaces (GUIs), and as such, focus on providing quantitative results of task-centric user performance. But unlike GUIs, new digital systems aim to augment and

coexist with users' real-world practices, naturally supporting much broader, informal and open-ended activities. As such, while traditional quantitative results of task-centric performance are valuable, they might not fully address and illustrate the wide range of benefits that contemporary interaction paradigms claim to offer.

This dissertation argues this limitation in HCI can be addressed by developing modern evaluation tools that appropriately quantify the broader impact of new interaction paradigms in users' activities. To do so it proposes to look at embodied cognition [S10], a new perspective in cognitive science that suggests thought processes are neither abstract nor centralized, but instead, rooted on multiple sensorimotor interactions with the physical environment [W02]. The use of these ideas in HCI is not new [e.g. A12], as they provide a compelling foundation from which to study the use of computer-augmented physical tools not only in task-centric activities, but also in real-world open-ended tasks. This will motivate the main contribution of this dissertation: the development of a modern evaluation tool that better accesses user performance and user sense making during interaction with contemporary digital systems. To do so, this dissertation focuses its work on two parallel topics: the design and development of novel computer systems based on a modern interaction paradigm that emphasizes multimodal, physical interaction – **tangible interaction**; and the identification, collection and use of an actionable set of metrics from embodied cognition – **epistemic actions**. Both these ideas will be presented in the following paragraphs.

Tangible interaction is an interaction paradigm that connects multiple physical forms to digital representations, allowing users to interact with data through their hands and bodies using physical objects (or *tokens*) as "interface, interaction object and interaction device" [HB06]. As with other modern interaction paradigms, proponents of the tangible paradigm have ascribed it

various benefits, including improved usability; increased levels of engagement and enjoyment; effectiveness is learning scenarios; and seamless support for collaboration [SH10]. A key factor underlying these prospective benefits is the integration of tangible interaction with the physical world, such that users are able to apply their naïve, but highly refined real world understandings of physical systems to interaction in the digital domain [JGH+08]. As such, by engaging users through bodily actions and physical manipulations, tangible interfaces are expected to facilitate what is known as *tangible thinking* [SH10] – using the environment to aid cognition in tasks at hand. An example of this type of action can be found in the work presented in LogJam's Video Logging system [CWP99], where users were observed structuring tokens around the application's surface in order to simplify and better organize the task. While this concept is compelling (and authors have argued it is poorly supported by purely virtual systems [SKM09]) there is still a lack of firm empirical evidence demonstrating the concrete advantages of designing user interfaces using physical, tangible technologies and approaches [A07, FGH+00, M07, G91, FTJ08].

Like with many other modern interaction paradigms, this lack is explained not only by a need for appropriate evaluation tools [SH10], but also because tangible interaction is still an emerging research area where most contributions focus either in innovative hardware and software solutions [e.g. LPI11] or on highly specific application areas such as the effects of tangibility on children's development and performance [MH09, XAM08, ZAM+12]. There are a few exceptions. Attempting to address more general questions, a piecemeal body of research has compared the effects of relying on tangible or graphical representations on user experience. While still providing valuable insights, this body of research has also attracted a range of methodological criticisms [XAM08, ZAM+12]. These include issues with the limited scope and generalizability of findings, as in Patten and Ishii's [PI00] comparison of the

use of space to organize information in graphical and tangible interfaces or in Marshall et al.'s [MCL10] discussion of the effects of using tangible versus graphical simulations of physical systems on adults' discovery learning. Other common issues include the fact that the advantages often credited to tangible interaction can be partly attributed to either novelty effects [ZAM+12] or to advantages inherent in enabling non-tangible technologies such as multi-touch displays (e.g. bi-manual input) [KST+09]. Finally, in comparisons between tangible and non-tangible systems, it can be challenging to ensure equivalence of the interfaces. For instance, in Soute et al.'s [SKM09] investigation of tangible and virtual game objects, the physical implementation introduced a range of novel functionality that was simply absent in the purely digital version. In sum, these results and observations make a strong case for the need of evaluation tools specific to novel interaction paradigms such as tangible interaction.

To generate such an evaluation tool, this dissertation will ground its work on embodied cognition [S10], the basis for concepts such as *tangible thinking*. The practical use of his this theoretical preposition in HCI is not new, and has been used to inform the design of various tangible systems. Examples include the MoSo Tangibles [BHA10], a set of interactive tokens that aims to teach sonic concepts such as pitch and tempo via the direct manipulation of embodied concepts such as speed, size and height; or NOOT [DRL+11], a tangible application that supports group brain-storming activities by characterizing and digitally augmenting external props (e.g., post-its, diagrams, sketches) as cognitive scaffolds. The use of embodied cognition for system evaluation in HCI is also not unprecedented. In his seminal paper, Kirsh et al. [KM94] introduced the concept of **epistemic actions** to better understand the difference in play strategies between skilled Tetris players and casual users. Kirsh described these actions as enabling users to structure and adjust the environment in order to have it contain or manipulate information

for them – "mak(ing) mental computation easier, faster, or more reliable". In sum, epistemic actions aim at altering the world so as to aid and augment cognitive processes (e.g., physically rotating a Tetris piece to see where it would fit, as opposed to mentally rotating it). They contrast with pragmatic actions that alter the world because some physical change is desirable for its own sake (e.g. fitting a Tetris piece in place). As such, this dissertation will expand on Kirsh's work by thoroughly exploring the use of epistemic actions as metric for user performance in modern HCI.

Research Questions

This dissertation will directly address four research questions:

1. Can Kirsh's findings [KM94] on Tetris players be replicated in other digital systems? Do epistemic actions increase user performance in those situations?

2. Do different input modalities (e.g., physical tokens, touch, mouse) to a common computer system evoke different epistemic actions? Kirsh introduced the concept by studying a purely graphical system (Tetris [KM94]), but literature on tangible interaction might suggest that physical representations are better equipped to afford these actions [e.g. SKM09].

3. If different input modalities do indeed evoke different epistemic actions, how do different design choices constraint or enable different epistemic behaviour?

4. And lastly, can epistemic actions serve as metrics in novel evaluation tools for modern interaction paradigms in HCI?

Dissertation Argument and Structure

As new interactive systems expand into our real-world surroundings, new evaluation tools are required to better understand the impact of their use. It is not sufficient to simply study how interfaces afford quicker acquisition of targets on a screen, or how they minimize human error [M92] – by existing in the real, physical world, new interfaces should be studied in order to capture the full range of benefits and full scope of the value that they offer their users. An example of such a novel interface paradigm is tangible interaction, an area within HCI that is concerned with the design and development of physical representations of digital data. This dissertation argues that a source for appropriate evaluation techniques for this field is found in embodied cognition, a perspective in cognitive science that explains how people make use of the physical tools around them to improve task performance and lower the cognitive demands of the problem at hand. To verify this claim, this dissertation will gather a set actionable metrics from embodied cognition and test their use as quantifiers of user performance across two comparative studies involving tangible systems. Once their value is empirically clear, a larger set of similar metrics will be collected and used to inform the development of a novel evaluation tool for tangible interaction. Finally, the validity, reliability and predictive power of this tool will be demonstrated through a user study where multiple researchers rely on the tool to measure not only user performance, but to provide a deeper understanding of user behaviour when dealing with a physical task. The remainder of this section illustrates how this dissertation is structured to address this argument.

The next chapter in this dissertation introduces the field of tangible interaction, describing a series of common applications and focusing on both the benefits credited to the field and the limitations that will be addressed

throughout this dissertation. The following chapter, chapter three, briefly describes the embodied cognition perspective, with a special focus on epistemic actions and how they can be used as metric in HCI. After these two introductory chapters, this dissertation introduces tiREC in chapter four, a tool for the systematic recording of quantitative data in tangible systems. All systems described throughout this dissertation are case studies that demonstrate the benefits of such a tool: easy setup and customization, and providing data that is coherent and rigorous across comparative studies of different interaction paradigms. The following chapter, chapter five directly addresses the first two research questions of this dissertation. It does so by measuring the impact of three epistemic actions in a repeated measures study of 36 participants. These competed against each other in an augmented game of Four-in-Row across three different interfaces – physical and graphical (touch and mouse input). The following chapter, chapter six, addresses the third research question of this dissertation. It does so by observing the different epistemic strategies employed by single and paired participants when interacting with different interface versions of tangible tabletop scheduling application. Finally, chapter seven, addresses both the initial and final research question presented. It does so by presenting a framework developed through a systematic literature review of 78 papers from various different fields of knowledge. The framework's value as metric for HCI is then put to a test in a 60 participant study across three different research groups. This dissertation ends by discussing the impact of its findings for the field of HCI as a whole, and presenting directions for future work.

Contributions

The outcome of this dissertation will include both empirical, design and tool oriented contributions. In terms of **empirical contributions**, this dissertation will demonstrate the importance of epistemic actions as quantitative metrics

for novel interaction paradigms such as tangible interaction. These metrics will then provide concrete empirical results attesting to the benefits of this interaction paradigm. These results will reflect the theoretical benefits that have been attributed to tangible interaction since Ishii and Ullmer's seminal paper on Tangible Bits [IU97]. This contribution ties directly with the **design contributions** of this dissertation, where design insights derived from each study are presented so that future tangible systems better support epistemic actions not only in task-centric efforts, but also in informal explorations of the task environment. Lastly, the **tool oriented contributions** of this dissertation come in the form of two separate tools for both developers and researchers looking to quantitatively evaluate their tangible systems. The first, tiREC, is a toolkit developed from the ground up to provide an open, coherent, and practical way to record the manipulation of physical and digital artifacts. The second, and the concluding contribution of this thesis, is the ATB framework. This is video-coding tool that categorizes 20 types of epistemic actions in three major action groups, enabling the fast and accurate measurement of these actions during interaction. The use of these tools will lead to important knowledge that can help guide and shape the design of interactive tangible systems in the future, opening the door to systems that not only improve user experience and user performance, but are less cognitively taxing and more natural to use.

TANGIBLE INTERACTION

2 Since the 1970s the dominant form of computer interface is based on the desktop computer, using mouse and keyboard to interact with a series of windows, icons, menus and pointers (WIMP). This kind of graphical user interface (GUI) serves as a general-purpose interface that emulates different tools and allows for many different interfaces to be represented digitally on-screen. As such, while the GUI paradigm excels at a task such as writing this dissertation, its bi-manual, generic input devices (mouse and keyboard) are shared across many other application domains, from productivity tools to games [KHT06]. Such constraints, combined with a newer understanding of the psychological and social aspects of human-computer interaction have led to an explosion of new post-WIMP interaction paradigms on a wide range of domains. These new interaction paradigms are becoming increasingly popular because through the use of novel input devices (e.g. gestural inputs, multi-touch surfaces) they leverage skills users normally employ in the real world [DBS09]. Furthermore, computers are also becoming embedded in everyday objects and environments, making computation more accessible, engaging, and responsive [R06].

One of such emerging post-WIMP interfaces is **tangible interaction**, which differs from GUIs by providing users with computers interfaces that rely on

well-defined physical forms that are specifically tailored to individual applications or systems [I08]. Such physical forms provide physical representations to digital information and controls, allowing users to literally interact with data through their hands and bodies. These interaction devices are implemented using a series of materials and technologies, allowing for physical objects to receive and interpret inputs (e.g. grabbing, squeezing, tapping, moving) and to provide a series of sensory stimulating outputs (e.g. adjustable textures and shapes, sound, visual effects). As such, these physical objects provide users with a parallel feedback loop, combining haptic information (passive or active) with digital feedback (visual or auditory) [UI00] – Figure 2.1. A tangible input device is then "simultaneously interface, interaction object and interaction device" [HB06].

The first instance of a tangible system can be traced back to the 1980s. Motivated by the fact that the CAD systems of that time were cumbersome and awkward to use, both Robert Aish [A79, AN84] and John Frazer's teams [F95, FFF80] looked for alternatives to input devices using physical models. As a result, they developed a computer that could create a digital model of a physical construction by scanning an assembly of blocks. The computer would take into account the location, orientation and type of each component, and could even provide suggestions on how users could improve on their design. Lastly, the final outcome could be printed by the user.

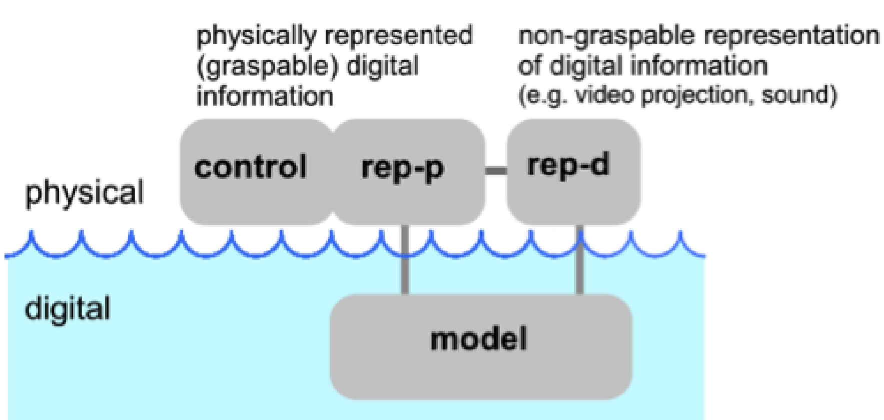

Figure 2.1. Interaction model for tangible interaction (MCRit model) [UI00].

But it was only in 1992 that Durrel Bishop created a seminal inspiration for tangible interaction: the Marble Answering Machine [A99]. In this device, colored marbles represented incoming or missed calls. As users place these marbles in different indentations of the system, they access different functions relating to the call represented by the marble: playing a voice message or calling the number back (see Figure 2.2). One year later, the issue "Back to the Real World" [WMG93] argued that the current desktop computers and virtual reality systems were very different from humans' natural environments. The article suggested that users should not be forced to enter a virtual world, but that the real world should be enriched and augmented with digital functionality. As such, computer interfaces should retain the richness of physical interaction by embedding computation in the users' tools, practices and environments [SH10].

It is easy to see how these ideas fueled not only tangible interaction, but several other HCI trends – from augmented reality to ubiquitous computing. However, it took several years for these ideas to take form. In 1995, Fitzmaurice et al. [FIB95] introduced the concept of a Graspable Interface, in which physical handles were used to manipulate digital objects. Two years later, Ishii and his students presented their work on Tangible Bits [IU97], which focused on mapping physical objects and surfaces with digital data – turning the real world in a computer interface (see Figure 2.3). While Ishii and his students worked on their vision for HCI, other research groups focused

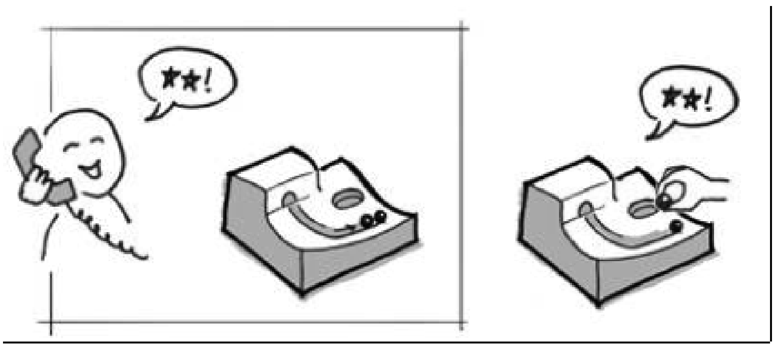

Figure 2.2. Durrel Bishop's Marble Answering Machine [A99]. A user picks a marble that represents a recorded call and places it in the indentation that instructs the system to play it.

applying these novel ideas to the design and development of applications for specific domains, digitally augmenting existing tools and artifacts. Examples include: the work of Wendy Mackay, with the use of flight strips in air traffic control and on augmented paper in video storyboarding [MF99]; the German Real Reality, which allowed for the simultaneous construction of physical and digital models [B93, BB96]; the work of Rauterberg and his team, which developed an augmented reality tabletop planning tool based on Fitzmaurice's graspable interface ideas – Build-IT [RFK+98]; the AlgoBlocks [SK93, SK95], created by Suzuki and Kato with the goal of supporting groups of children in learning how to program; or Cohen et al.'s Logjam [CWP99], a tangible system that supported video coding and logging. Since then, and with help from the fields of ethnography, situated and embodied cognition, and phenomenology [W93b], tangible interaction continues to develop, striving for a seamless connection between the digital and the physical world [SH10].

Terminology and Classifications

Tangible interaction encompasses several post-WIMP ideas that rely on different concepts, approaches and technologies. Some of these ideas were born from the work of Ishii on Tangible Bits [IU97], while others from particular types of tangible systems. These include: ambient displays [e.g., W93], tangible augmented reality [e.g., BKP01, LNB+04]; tangible tabletop interaction [e.g., JGA+07]; embodied user interfaces [e.g., FKY05, FGH+00]; or tangible

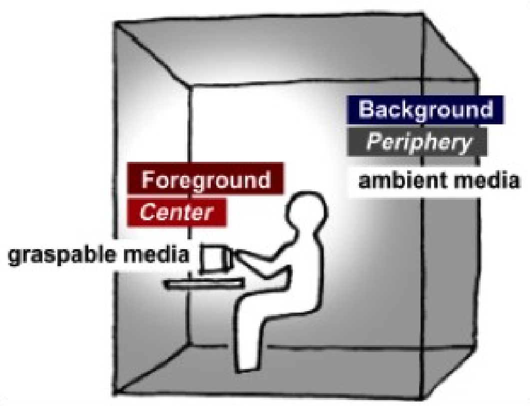

Figure 2.3 Center and periphery of the user's attention within a physical space [IU97].

autonomous interfaces [e.g., NK14]) – see Figure 2.4.

Additionally, several other authors have provided classifications for modern interaction paradigms that fall into what this dissertation considers to be tangible interaction. The first is **tangible computing**, a term that arose from multiple concepts presented by Dourish [D01] as being fundamental to the goal of integrating computation into our everyday lives. These concepts included tangible systems, ubiquitous computation, augmented reality, reactive rooms and context-aware devices. The purpose of tangible computing is to distribute computation over many specialized and networked devices in the environment, augmenting the real world so it can accurately react to the users' actions. This interaction should not focus on solely one input device, but combine a series of objects and tools. There also should not be an enforced sequence of actions that users must follow, and the design of the tangible objects should inform the users of their purpose and functionality. According to Dourish's classification, tangible interfaces differ from the other HCI approaches as it allows for computational representations to be artefacts in every sense of the word, allowing users to directly act upon, lift up, rearrange, sort and manipulate them in the physical world. Additionally, manipulating such objects can accomplish several different non-task-centric goals – e.g. moving a prism token in Illuminating Light [UI98] can be done simply to make space, to explore the system's output, to use it as a tool or to explore the entire system as a tool. The users can freely and effortlessly change their attention between all these different goals because computation is embodied.

Figure 2.4. From left to right: sketch examples of an ambient display, a tangible augmented reality object, a tangible tabletop application and an embodied user interface [SH10].

The second classification to fall under tangible interaction is **reality-based interaction**. Proposed by Jacob et al. [JGH+08], it is a unifying framework for various emerging interaction styles in HCI. These styles include virtual and augmented reality, ubiquitous and pervasive computing, and physical interaction. The common denominator between them is the fact that they all aim to take advantage of users' real-world skills, trying to narrow the gap between engaging with digital and physical tools. This framework is comprised of four main themes of interaction that are naturally leveraged in the real world (see Figure 2.5):

Naïve Physics: the subconscious human understanding of physical principles such as gravity, friction, velocity, the persistence of objects and relative scale.

Body Awareness and Skills: the familiarity humans have with their own bodies (e.g. coordination, range, relative position of limbs), independent of their surroundings.

Environment Awareness and Skills: the sense of physical presence humans have in a physical environment. The natural clues available facilitate humans' sense of orientation and spatial understanding.

Social Awareness and Skills: the awareness that other people share the same environment and the skills necessary to interact with them. These include

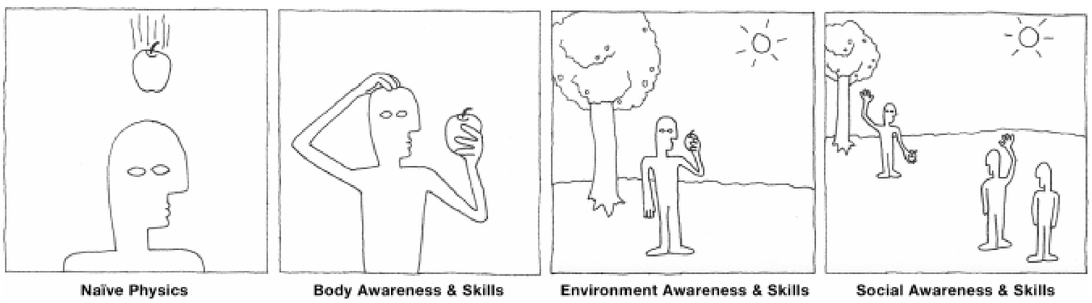

Figure 2.5. The four central themes in Jacob et al.'s [JGH+08] reality-based interaction.

verbal and non-verbal communication, the ability to exchange physical objects and the ability to collaborate with others on a task.

To date, most tangible designs and rely on the users' understanding of **naïve physics** (e.g. grasping and manipulating tangible objects) and of **social awareness and skills** (e.g. sharing of interaction space and tangible objects). Additionally, Jacob et al. suggests this trend towards reality-based interaction may reduce the mental effort required to operate current computer systems. As such, the aim of emerging interaction styles in HCI should be to reduce the gap between the users' goals and the means to execute them. Jacob et al. believes interaction designers should only give up grounding their work on the physical world in return for other desired qualities, such as: **expressive power** (the ability to perform a variety of tasks within the application domain), **efficiency** (the ability to quickly perform a task), **versatility** (the ability to perform a variety of tasks from different application domains), **ergonomics** (the ability to perform a task without fatigue or the risk physical injury), **accessibility** (the ability to perform a task in the presence of physical limitations) or **practicality** (the ability to develop systems that are practical to develop and produce). Additionally, the descriptive nature of Jacob et al.'s framework enables developers to: analyse and compare alternative concepts and designs; bridge gaps between tangible systems and other research areas; and apply lessons learned from developing in other interaction styles. In sum, the reality-based interaction framework strives to popularise the development of tangible interaction by creating richer interactions using the users' natural skills and the knowledge of the physical world.

The third and last classification to fall under the umbrella the **tangible interaction** is a framework of the same name. Proposed by Hornecker and Buur [HB06], it is used to describe the group of fields that relate to, but are broader than simple tangible interfaces. The authors give examples of systems

developed by artists and designers, which have very rich physical interactions, and thus, share some of the characteristics of traditional tangible applications. These systems tend to enable rich full body interactions [BJD04], focusing more on the expressiveness and meaning of body movements, and less on the tangible objects and the digital information being manipulated. This framework classifies tangible interaction in three different categories:

Data-centered view: the most traditional and computer-oriented of these views, where tangible objects represent and manipulate digital information.

Expressive-movement view: this view focuses on bodily movement, rich expression and physical skill – meaning is created in the interaction.

Space-centered view: more relevant in the arts, this view emphasizes on interactive and reactive spaces, combining tangibles objects with digital displays or sound installations. Full-body interaction – the use of the body as both the interaction device and display – is also characteristic of this view.

In sum, tangible interaction is the terminology preferred by the HCI design community due to its broader focus on the user experience [BG10, HFA+07]. It encompasses tangibility, physical embodiment of data, full-body interaction and systems that are embedded in real spaces and contexts. Nevertheless, in this dissertation tangible interaction will refer not only to this latter classification, but to all uses of the term discussed in this section.

Application Domains

To better contextualise the essence of the work being pursued in tangible interaction, this section will introduce the most popular application domains for tangible systems.

Figure 2.6. A Topobo [RPI04] assembly. Each physical part can be individually programmed by simple demonstration.

Child Development: Researchers such as Bruner and the late Piaget emphasized that embodiment, physical movement and multimodal interaction are of paramount importance to the development of children [A07, OF04]. Other researchers have also argued that gesturing can support thinking and leaning [G05], and that physical interaction can help children plan and reflect on their activities [M07]. Both Antle [A07] and Zuckerman [ZAR05] have applied this knowledge to the development of frameworks that provide design guidelines for tangible systems. Examples of systems that reflect these ideas include: the Lego Mindstorms, born from the work conducted at the MIT Media Lab Lifelong Kindergarten group [R93]; Smart Blocks, which allows children to explore physical concepts such as volume and areas [GSH+07]; and Topobo [RPI04], a tangible system that enables children to learn about balance, movement patterns and anatomy through the assembly of a physical robot whose parts can be programmed individually though demonstration (see Figure 2.6). Additionally, researchers have argued that tangible systems are particularly suitable for children with special needs [VSK08], such as speech impairment [HHO08, HHO09]. Tangible objects slow down interaction, provide a rich sensorial experience, support collaboration and allow children to train their perceptual-motor skills while retaining control of the interaction [SH10]. Finally, several tangible systems have been developed as diagnostic tools that measure children's cognitive level, spatial abilities and overall development [SIW+02, WKS08]. This is normally achieved by recording the interaction and analysing the steps and mistakes performed during play (e.g. building physical structures).

Information Visualization: It is very challenging to physically alter tangible objects. As such, these normally rely on rich multimodal representations that convey haptic, visual and auditory feedback to portray different digital information. Examples of tangible systems for information visualization include: an application for 3D neurosurgical visualization that enables surgeons to simulate surgical slices by holding a flat plate close to a doll's head [HPG+94]; GeoTUI [CRR08], a system that allows geophysicists to use tangible props to cut planes on geographical maps projected on a physical surface; or Ullmer et al.'s [UIJ05] famous tangible query interface prototypes, which use tangible objects as database parameters such as queries, views or Boolean operands. In addition to these examples, several others are finding their way into home and office settings [e.g. EB09, HIW08].

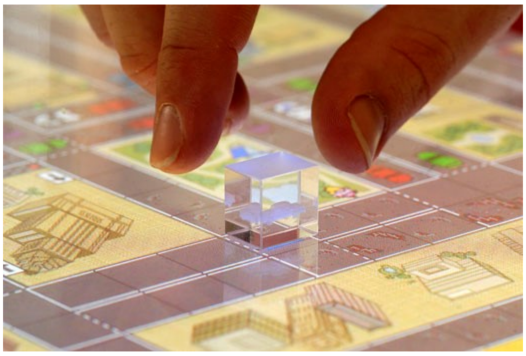

Figure 2.7. The Philips EnterTaible project [LBB+07].

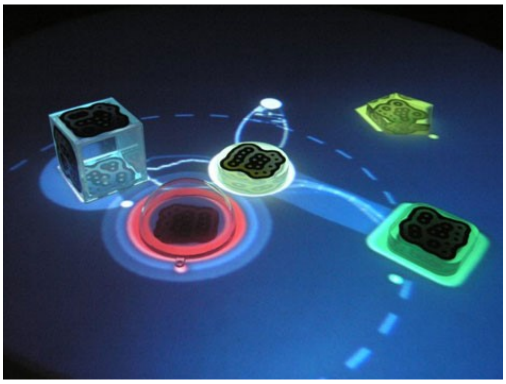

Figure 2.8. The reacTable [JGA+07], an extremely popular tangible music system.

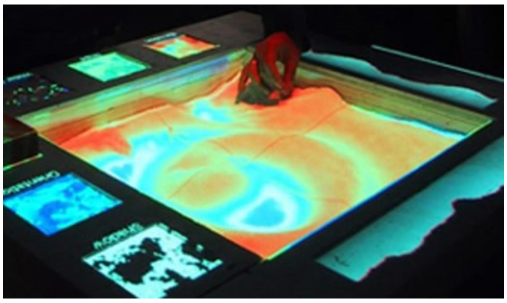

Figure 2.9. The SandScape [I08b], a tangible system that allows users to interact with physical sand to manipulate a digital landscape in real-time.

Tangible Programming: Suzuki and Kato coined the term "tangible programming" in 1993, using it to describe their AlgoBlocks system [SK93, SK95]. The goal of this application was to teach children how to program through a tangible game where the connection of physical blocks representing the Logo programming language would guide a digital submarine. Since then

several tangible systems for programming have been developed. These normally fall in one of two groups: **programming by demonstration** [CH93] (or programming by rehearsal [L93]), where children teach electronic robots to move by demonstrating a set of motions or gestures; and **constructive assemblies** [UIJ05], where children connect modular pieces to create a tangible structure that represents a simple programming syntax. Examples of former include Topobo [RPI04] (see Figure 2.6), Curlybot [FMI00] and StoryKits [SDM+01]. Examples of the latter include AlgoBlocks [SK93, SK95], Digital Construction Sets [M04], Electronic Blocks [WP02], GameBlocks [S07] and Tern [HSJ08]. All these examples are aimed at children, and as such could be understood as systems for **child development**. The truth is that tangible programming has no age restriction, as the researchers at the Mads Clausen Institute in Denmark demonstrate us [SCB07]. Their work uses tangible interaction in the context of industrial work, supporting the configuration work of service technicians. This work is especially important because it makes an attempt to bring back some of the advantages that tangible interaction shares with traditional mechanical interfaces: motor memory, real-world skills and visibility of action.

Entertainment: While not a traditional tangible interface, the phenomenal success of the Nintendo Wii shows the market potential of embodied, tangible systems for entertainment. This potential has been explored by several museums, who feature installations that use tangible interaction to engage with visitors. A classic example is Joseph Paradiso's Brain Opera installation at the Vienna Haus der Music (Museum of Sound), where a room full of tangible objects generates sound in response to visitors' movement, touch and voice [P99]. Other examples of tangible systems for entertainment include augmented board games that combine the social atmosphere of traditional games with digital gaming experiences (e.g., the Philips EnterTaible project [LBB+07] – see Figure 2.7); Leitner et al.'s [LHY+08] mixed reality gaming

table, a system that tracks physical objects to transform them into path or obstacles in a projected virtual car race; or IOBrush [RMI04], a drawing tool that allows children to explore colour and texture via a tangible paintbrush. Finally, Sturm et al. [SBG+08] identified key design aspects that future tangible systems for entertainment should encourage: support for social interaction; simplicity but adequate challenge and goals; and a motivating feedback system.

Music and Performance: Music applications are among the most popular areas of research in tangible interaction. They usually target two distinct groups of users: the novices, who want intuitive and carefree interaction; and the professionals, who look for expressiveness, legibility and visibility when performing for a crowd [KHT06]. Several commercial systems are already available, and include: the AudioCubes [SV08], a tangible system comprised of physical cubes that can detect and communicate with each other to generate sound; the Block Jam [NNG03], a dynamic sequencer built from physical blocks that can be attached to one another; and mixiTUI [PH09], a tangible sequencer that physically represents music and sounds clips, and enables several musical operations such as loops and effects. Kaltenbrunner [K09] classified these and several other musical systems into five categories: the systems that **embed** music and sound within physical artefacts which can be rubbed, squeezed or moved (e.g. the Squeezables [WG01]); **musical building blocks** that rely on the spatial properties of individual or groups of physical blocks to repeatedly generate or manipulate sound (through e.g., stacking, attachments or by simply having blocks close to one another); **token-based sequencers**, who are laid on a surface and continuously scanned to produce sound depending on their spatial position and their physical properties (e.g., colour); **interactive touch-based surfaces** in which music is produced as a result of interactions with tangible objects (e.g., AudioPad [PRI02], reacTable [JGA+07] – see Figure 2.8); and finally, simple commercial

toys that generate music by detecting the presence and sequence of tangible objects in specific physical slots (e.g. Neurosmith's MusicBlocks, Fisher-Price's play zone music table). Finally, Jordà [J08] has identified several desired properties for tangible musical applications: natural support for collaboration and sharing of control; uninterrupted, real-time interaction with multidimensional data; and support of complex, skilled, expressive, and explorative interactions.

Problem-solving: The most iconic tangible systems have been problem-solving applications. Because the external representation of a problem is known to profoundly impact the strategies humans employ to solve it, it is easy to perceive tangible applications as having a deep impact in user performance [VP05]. Additionally, the existence of physical constraints in a tangible interface creates affordances (e.g. racks, slots) that communicate the system's syntax. This decreases the learning time and consequently lowers the threshold for participation [UIJ05]. As such, these tentative benefits of tangible interaction have naturally led to the development of a multitude of problem-solving applications. Examples include: Urp [UI99], a system that supports urban planning through the manipulation of physical objects that simulate buildings' shadows, proximity or wind paths; SandScape and Illuminating Clay [I08b, PRI02], two systems where users are free to explore a digital landscape by deforming sand and clay with their hands (see Figure 2.9); and the MouseHaus Table [HYG03], a system that allows users to collaboratively conduct a pedestrian simulation through the placement and movement of everyday objects upon a physical surface. Examples of more abstract systems also exist. These include: the Senseboard [JIP+02], a system that enables users to organize and manipulate abstract pieces of information by grouping physical pucks on a vertical grid; Pico, a system that manages different application variables by relying on a tabletop surface to sense and

adjust physical objects autonomously; and finally TinkerSheets [ZJL+09], a system that provides a simulation environment for warehouse logistics.

The work developed for this dissertation reflects this diversity of application domains. The systems to be introduced in the following chapters can be characterized as applications for: Music and Performance (chapter four), Entertainment (chapter five), Problem-solving (chapters four, five and six) and Information Visualization (chapter six). As such, the findings presented in this dissertation should be applicable to a wide range of future work in this field.

Technologies and Tools

Interaction with traditional GUIs normally relies on a core set of input devices (mouse and keyboard) and output devices (screen and speakers). This is not the case for tangible systems, as developers employ a wide range of existing and custom-made technologies to bring their solutions to life. This process is made harder by a lack of public standard solutions for development and because tangible systems are likely required to track several users and physical objects in an open-ended, physical environment. Following is a collection of the most commonly used technologies in tangible interaction, as well as the most popular tools available.

Radio-Frequency Identification (RFID) is a wireless radio-based technology that is able to detect tagged objects when in range of a tag reader. Normally, an RFID tag contains an integrated circuit for storing and processing information and an antenna for receiving and transmitting data. The communication between a tag and a reader only occurs when both are in proximity (up to 100 meters). Once detected, the tag reader extracts the ID of the tag and delivers it to the appropriate tangible application. Examples of

tangible systems that use RFID technology are the MediaBlocks [UI99b] and the Senseboard [JIP+02]. Both rely on this technology to individually tag physical blocks and map them to specific digital data.

Figure 2.10. A fiducial marker. A popular tool to provide cheap and accurate tracking of physical objects with computer vision.

Computer Vision: A popular technology within tangible interaction, computer vision enables applications to detect the position, orientation, colour, size and shape of multiple objects in real time. The cheapest, most accurate way of achieving this is by individually tagging physical objects with visual tags (e.g., fiducial markers – see Figure 2.10). Tangible systems that employ computer vision are typically composed of a high-quality camera (that scans the task environment), a light-weight LCD projector (that provides real-time graphical output) and a software package that is responsible for identifying tangible objects. Examples of such systems include the Urp [UI99], an application for urban planning; and the reacTable [JGA+07], a platform for musical performance. There are also several examples of software libraries that support the development of tangible systems based on computer vision: the ARToolKit [KBP+00, KB99] and the reacTIVision [JGA+07], which support the tracking of fiducial markers; and Papier-Mâché [KLL+04], which is capable of detecting electronic tags and barcodes. Finally, with the introduction of Microsoft's Kinect sensor and software library [Z12] it is now practical and easy to track multiple users and individual parts of their bodies.

Figure 2.11. An Arduino UNO microcontroller [B09].

Microcontrollers, Sensors, and Actuators: Microcontrollers are small and inexpensive computers that can be embedded in physical objects and in the

environment itself. They receive information from their surroundings through a series of sensors and respond accordingly through different actuators. Their sensors are capable of detecting light intensity, reflection, noise level, motion, acceleration, location, proximity, position, touch, altitude, direction, temperature, gas concentration and even radiation. Actuators on the other hand, are capable of producing light (e.g., through LEDs), sound, motion (e.g., through motors and electromagnets) or haptic feedback. Examples of tangible systems that employ microcontrollers include the Senspectra [LPI07], a physical modelling toolkit for analysing structural strain; Pico [PI07], an interactive system that uses actuation to move tangible objects over its surface; or Navigational Blocks [CDJ+02], a tangible application for navigating and retrieving historical information using haptic feedback. Examples of development environments for these tangible systems include the Arduino [B09], the Handy Board [M00], the Handy Cricket [MMS00] and Lego Mindstorms NXT [B03].

Phidgets [GF01] is a tool that consists of *plug and play* devices (e.g. sensors, boards, actuators, I/O) that support developers in the implementation of tangible interaction prototypes. One of the advantages of developing with Phidgets is that all devices are centrally controlled by a computer rather than through multiple microprocessors embedded in the tangible objects. Additionally, Phidgets offers a software API for a variety of development environments. iStuff [BRS+03] is very similar to Phidgets in concept, but relies instead in set of wireless devices controlled in a Java environment that defines high-level events and dynamically maps them to input and output events. Exemplar [HAM+07] is

Figure 2.12. Three Siftables devices [MKM07]. Each device is capable of sensing physical input and providing graphical output to the user.

another toolkit similar to Phidgets in the way it leverages central control through a computer. With Exemplar, developers enact sensor-based interactions with their systems (e.g. shaking an accelerometer) and iteratively refine a desired input pattern.

Arduino [B09] is a toolkit consisting of a special board and programming environment (see Figure 2.11). The Arduino board differs from Phidgets by interfacing with standard electronics parts, requiring developers to physically wire, circuit build and solder. Due to the popularity of this toolkit, numerous compatible spinoffs have been produced, such as the Seeeduino Film [S14] (which offers a bendable board) or the Digispark [O13] (which offers a coin-sized board).

Siftables [MKM07] is a platform that relies on individual devices that sense physical input and provide graphical output to the user. These devices are capable of sensing and communicating with each other, and as with Phidgets, they are also controlled by a central computer (see Figure 2.12). Unlike Phidgets or Arduino, which enable developers to create their own physical devices by combining sensors and actuators, the Siftables platform enables a rapid (if constrained) prototyping experience out-of-the-box by constraining developers to the (broad) capabilities of each device.

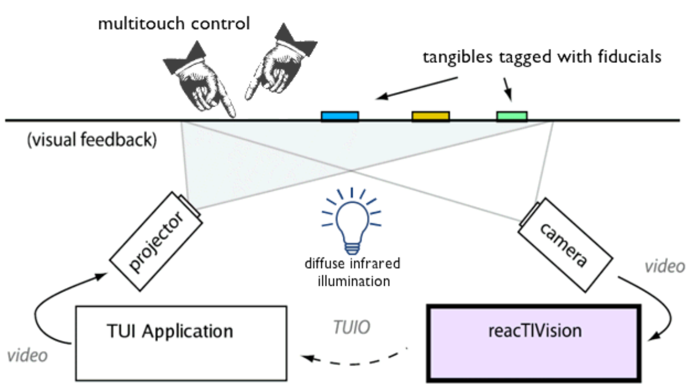

Figure 2.13. The reacTIVision framework diagram [JGA+07].

ARToolKit: As mentioned before, the ARToolKit [KB99, KB00] is a computer vision marker tracking library that allows developers to create augmented reality, and consequently tangible applications. In addition to tracking the position and orientation of visual markers, ARToolKit also allows computer graphics to be drawn directly over the markers.

reacTIVision: The reacTIVision [JGA+07] is a computer vision framework created for the development of tangible systems that combine multi-touch input with the tracking of tangible objects tagged with fiducial markers (see Figure 2.13). One of the advantages of the reacTIVision framework over other toolkits is its distributed architecture, separating the sensing and tracking from the user applications.

Benefits

As with other modern interaction styles, tangible interaction has been associated with a series of benefits for user performance and experience. These include seamless support for **collaboration** [HSC+09, KHT06, SH10, GSO05, FT06, FTJ08, FTJ08b], as tangible systems naturally support multiple access points that afford simultaneous interaction, group awareness and easier participation and coordination; improved **usability** [F96, N88, B00, WDO04], as tangible objects enable parallel space-multiplex input that is function specific; increased levels of **engagement** and **enjoyment**; effectiveness in **learning** scenarios [e.g., A12]; and being **powerful** Ubicomp devices [WMG93] that inhabit the same world as we do. Another benefit of tangible interaction is what is called *tangible thinking*. When compared to purely graphical interfaces, tangible systems are expected to better leverage the connection of body and mind by facilitating mental processes through bodily actions, physical manipulation and tangible representations [ZN94]. This is because our understanding of the world is shaped by both our

physical bodies and how they interact with the world and the objects that inhabit it [KHT06, T07] – it is locomotive experiences that allow infants to develop their spatial cognitive skills, and it is through bodily interactions with physical objects that children learn abstract concepts [RMB+98]. As such, the physical objects that populate tangible interaction can be understood as *thinking props* that allow the offloading of cognition and memory [HM11]. These physical tasks representations can also allow users to perceptually infer what they can do with a system (and how), decreasing the need for explicit rules [Z97, ZN94]. In addition to providing users with a physical interface that can aid cognition, tangible interaction also enables the use of unconstrained gestures during interacting (as the users' hands are not confined to a keyboard and mouse). This can further lighten the cognitive load during task performance [AKY00]. Additionally, tangible systems that rely on gestures for input are expected to take advantage of users' kinesthetic memory [S99].

While these benefits are very encouraging – and it is true that tangible interaction can leverage users' accrued knowledge of interacting in real world [H12] – there is still a lack of empirical, quantitative data demonstrating the benefits of tangible systems aiding user performance and lowering cognitive load. A growing body of work has begun addressing this issue by empirically comparing tangible interaction with other interaction paradigms. While these efforts are commendable, they have also attracted a range of methodological criticisms [XAM08, ZAM+12]. These include issues with the limited scope or generalizability of findings [e.g., PI00, MCL10]; the fact that some of the advantages credited to tangible interaction can be at least partly attributed to either novelty effects [ZAM+12] or to advantages inherent in enabling non-tangible technologies such as multi-touch displays (e.g. bi-manual input, unconstrained gestures) [KST+09]; and finally, in comparisons between tangible and non-tangible systems it can be challenging to ensure equivalence

of the interfaces – in several instances the tangible representation introduces a range of functionality that is simply absent in the purely digital representation [e.g. SKM09].

Challenges and Limitations

While tangible interaction is an exciting and promising field, it also presents developers with a series of challenges that stem from both the novelty and the natural qualities of the paradigm. Unlike digital objects that are malleable, easy to create, modify, replicate and distribute, tangible objects are normally rigid and static [PNO07]. Many issues arise from this fact, such as a difficulty in correcting errors (there is no undo function in the physical world) or in the automatic replication of actions [KST+09]. These specific limitations are the origin of one of the most recent areas of research within tangible interaction – tangible autonomous interfaces composed of objects capable of independent movement or motion [NK14]. Physical objects have also an additional undesirable characteristic: they take up physical space. As such, tangible systems can prove very hard to scale up to encompass large numbers of objects.

Another big challenged for the field is how to design and develop tangible systems. Much of the established body of work in HCI, which has largely focused on desktop GUIs, is simply not applicable to this new paradigm [A03]. Reflecting this lack, frameworks stipulating how to structure and develop tangible systems have been developed over the years [MH09]. Their goal is to provide general guidance and direction – e.g., in [HB06] the authors describe a framework of four themes for design of tangible experiences (tangible manipulation, spatial interaction, embodied facilitation and expressive representation). While valuable, these frameworks provide little actionable information regarding how tangible systems can be designed to

maximize usability, to reflect the strengths and weaknesses of the human mind, or to be "cognitively ergonomic" – central characteristics of physical interfaces. This lack of concrete and practical design knowledge is derived from three distinct challenges:

1. The body of work on tangible interaction has led us to believe that tangible systems can leverage both the positive aspects of physical inputs and users' prior knowledge of the world without much consideration for the design rationale. While naturalness and intuitiveness [DWF+04] were the biggest 'selling points' when tangible interaction was first introduced , it is now becoming clear that much more effort is needed to fulfil that promise [H12].

2. It is very challenging to document best practices when prototyping tangible systems. As mentioned earlier, unlike GUIs, tangible interaction has almost a limitless number of input and output possibilities. As such, each time a new tangible system is prototyped researchers are most likely required to learn a new hardware setup or software toolkit. Although toolkit programming substantially reduces the time and effort required to build functioning systems, they normally fall short of providing a comprehensive set of abstractions for specifying, discussing, and programming tangible interaction within an interdisciplinary development team [SH10]. This issue is illustrated in the tangible systems developed for this dissertation: some systems relied on computer vision and the reacTIVision framework [JGA+07] (chapters four and six), others relied in combining microcontrollers with several sensors and actuators (chapter five).

3. And lastly, there has been limited progress in the development of specific evaluation tools for tangible interaction. This is not to say that tangible systems are rarely evaluated – numerous studies provide detailed and appropriate information on participant feedback in both lab [e.g., BHA10] and

in-the-wild settings [e.g., GBB+04]. While most of these examples support the notion that tangible systems have a special impact on users' sense making and thought processes [S10], there is little work providing empirical evidence that tangible interaction paradigms are more effective that traditional HCI interfaces.

The remainder of this dissertation will be dedicated to addressing this latter challenge of tangible interaction. Most evaluations of tangible interfaces are formative in character – they seek to isolate appropriate characteristics to inform system design. Such studies are rarely critical of the design rationale or of the fundamental value of tangibility [ZAM+12]. As such, what tangible interaction is missing are paradigm-specific tools that provide quantitative data of user performance across either different interaction paradigms or different tangible systems and designs. This lack has led to difficulties in robustly demonstrating the value of incorporating tangible interaction techniques into a system design. While past work has compared tangible interfaces to other forms of interaction, such as keyboard/mouse or multi-touch input (e.g. [RGB+10]), these studies have been hard to generalize or required the users to perform tasks that were highly simplistic [TKI10]. Furthermore, they often describe how participants simply reported they felt more satisfied using tangible interaction [e.g. RGB+10] or that they found it easier to control and manipulate digital information through physical props [e.g. TKI10]. Such studies are typically well documented, and the findings are solid and replicable. However, it is clear that such qualitative data can only support very general claims and has limited applicability. Finally, this dissertation believes that providing a solution for this challenge will indirectly address both challenges one and two.

EMBODIED COGNITION, EPISTEMIC ACTIONS

3 In order to understand the popularity and excitement around tangible interaction, one must first understand the ways in which physical embodiment and physical objects can influence, effect and constrain our cognitive activity. To do so, this chapter presents a brief literature review of the body of work pertaining to embodied cognition, a perspective in cognitive science that grants the body and its interactions with the world a central role in how the mind works.

Traditionally, researchers have viewed the mind as an abstract information processor [CMN86], regarding connections to the outside world as of little theoretical importance for human cognition. Perceptual and motor systems were thought to serve merely as peripheral input and output devices, not relevant to understanding *central* cognitive processes [W02]. It was only in 1980's that researchers from disparate fields started to argue for a broader understanding of cognition: linguists argued that abstract concepts were based on metaphors for bodily and physical concepts [LJ80], and researchers in behaviour-based robotics began to develop artificial intelligence that was based on routines of physical interaction and exploration, rather than internal representations for abstract computational thought [B96]. The work on embodied cognition arose from these approaches and has recently attained

high visibility. Proponents of the field argue we evolved from animals whose neural resources were devoted primarily to perceptual and motoric processing, and whose cognitive activity consisted largely of interactions with the environment. As such, they argue human cognition is neither centralized nor abstract; it is instead rooted in processes of sensorimotor activity [W02]: as a person moves through an environment, his or her locomotion will produce opportunities for new perceptions while at the same time erasing old ones; in turn, the new perceived features and affordances will reveal opportunities for new activities. This creates a feedback loop where motion influences perception, which influences future motion, which in turn determines new perceptions, and so on [S10].

Embodied cognition further suggests that intelligence is intrinsically supported by physical tools provided by our social and cultural world, and it is through repeated interaction with the surrounding environment that people create structures that advance and simplify their cognitive tasks [A03]. This means that people "tend to recruit, on the spot, whatever mix of resources will yield an acceptable result with a minimum effort" [C08]. These ideas also help explain how humans store and access memory. Glenberg [G97] argued that traditional accounts of memory focus too much on the passive storage of information. He defends the idea that patterns stored in memory reflect the nature of bodily actions and their ability to mesh with different situations during goal pursuit. As such, perception of relevant objects triggers affordances for action stored in memory. Conversely, reasoning about future actions relies on remembering affordances while suppressing perception of the environment. Simulation also appears central to constructing future events based on memories of past events [SA07]. When people view a static configuration of gears, for example, they use simulation to infer the direction in which a particular gear will turn. Numerous sources support the use of simulation in these tasks [e.g. B08]. The time to draw an

inference is often correlated with the duration of a physical event, such as how long a gear takes to turn. It has also been shown that carrying out associated actions (e.g. moving your hand like a gear) can improve inference [S99b]. All these ideas have helped define embodied cognition under four perspectives:

Cognition is situated: Situated cognition relates to cognitive activity that takes place in the context of a real-world task, and inherently involves perception and action. Examples include driving or holding a conversation – interacting with the things that the cognitive activity is about [W02].

Humans offload cognition onto the task environment: If cognition is influenced by the task environment, physically altering the environment has direct impact in the cognitive load of the task. Humans exploit this on a daily basis, making the environment hold or even manipulate information for them, and harvesting that information on a need-to-know basis. Clark classified this as external *scaffolding* - how the brain offloads some of its cognitive duties onto the environment, instead of doing all of the computational work on its own [C01]. An example of this mental shortcut is how people look for Kodak equipment in a store, focusing on finding a distinct yellow package instead of a specific textual element [C97]. Physical landmarks are also an example of external cognition in cities. By offloading mental effort onto complex urban environments, the residents and visitors of cities are able navigate without experiencing mental overload [M06]. A particular kind of scaffolding is known as **epistemic actions**. These actions aim at altering the world so as to aid and augment cognitive processes (e.g. rotating a Tetris piece to see where to fit it [KM94]). They contrast with *pragmatic actions* which alter the world because some physical change is desirable to reach a task-specific goal (e.g. fitting a Tetris piece in place [KM94, CC98]). Additionally, people perform epistemic actions to either reduce: **space complexity** (memory involved in mental computation); **time complexity** (number of steps involved in mental

computation); or to reduce **unreliability** (probability of error of mental computation) [KM94].

Another recently proposed theory that explains how humans offload cognition into the environment is what Kirsh calls **projection** [K09b]. According to his work, projection is a key part of a cyclical process of solving problems in the world, in which users act, observe the result of their actions and consider their next move [K09c]. In this cycle, projection sits between *perception*, which refers to what is sensed from the real world, and *imagination*, which refers to entirely mental constructs. Between these poles, projection refers to mental augmentations of reality that are anchored and grounded on perceived external structures [K09b]. For example, projection is the process that occurs when a person looks at a piece on a chessboard and is able to visualize the possible (or even the good vs bad) moves. In this framing, when people finally act, they externalize a structure that was initially mental. As such, in projection an epistemic action can serve two crucial purposes. One is lowering the cognitive cost of projecting by instantiating some of the mental constructs in the real world (e.g. lifting a piece from the chessboard to better understand the impact of moving it). A second is to nurture additional projections, e.g. hovering a chess piece over a possible future location on the board to better envision future moves [K09b]. The potential of epistemic actions for HCI and tangible interaction will be discussed in more detail in the following section.

Cognition is for action: The function of the mind is to guide action, and cognitive mechanisms such as perception and memory must be understood in terms of their contribution to task-appropriate behaviour [W02]. Glenberg [G97] argues that the traditional approach to memory as “for memorizing” needs to be replaced by a view of memory as “the encoding of patterns of possible physical interaction with a three-dimensional world”. This approach

shows that we conceptualize objects and situations in terms of their functional relevance to us, rather than neutrally or "as they really are" [W02].

Offline cognition is body-based: Even when decoupled from the environment, the activity of the mind is grounded on mechanisms that evolved from interaction with the physical world: mechanisms of sensory processing and motor control. Mental imagery (e.g., auditory and kinaesthetic imagery) is an obvious example of the mind simulating external events [PFD+95]; while the working memory appears to be an example of a symbolic offloading where information is offloaded onto perceptual and motor control systems in the brain. In both these examples, rather than the mind operating to serve the body, we find the body (or its control systems) serving the mind [W02]. As such, areas of human cognition previously thought to be highly abstract now appear to be yielding to an embodied cognition approach. It appears that offline embodied cognition is a widespread phenomenon in the human mind [W02].

Embodied Cognition in HCI

The use of embodied cognition in HCI is not new: in his seminal book [D01], Dourish framed key ideas from embodied cognition in a HCI perspective. This work has inspired later researchers to follow a similar approach in the field of tangible interaction. For example, Klemmer et al. [KHT06] developed five themes to guide interaction design that are deeply rooted in principles from embodied cognition, especially the themes *thinking through doing*, *performance* and *visibility*. Hurtienne et al. [HIW08] illustrated several examples of user centered design that are based on *image schemas*, a specific concept from embodied cognition that describes how different abstract representations are rooted on core bodily actions. And Antle et al. [ACD09] outlined the benefits of using embodied metaphors [LJ80b] in the

development of an interactive environment. All these theoretical understandings are of great practical value and have been used to inform the design of a number of tangible systems. An example is MoSo Tangibles [BHA10], a set of tangible artifacts that teach children abstract sound concepts through the movements and gestures they support. Three of these artifacts are illustrated in Figure 3.1, and demonstrate how *pitch* was taught to children by relying on three different image schemas: low-high, near-far and slow-fast. Another example is the NOOT [DRL+11], a tangible system that facilitates the sharing of moments of reflection during brainstorming activities. The design of the prototype was strongly motivated by a single property of embodied cognition: how users *offload* taxing mental processes (e.g., remembering) through the physical manipulation of the environment. These examples demonstrate how different theoretical constructs from embodied cognition can successfully motivate and guide the design and development of modern interactive systems.

This dissertation believes these same constructs can be applied to the systematic and quantitative evaluation of tangible systems. The potential of embodied cognition for evaluation in HCI was first demonstrated by Kirsh et al. [KM94], as they introduced a novel metric for user performance that was explicitly non-goal directed – epistemic actions. Two decades later, Antle et al. [AW13] expanded this idea to an actionable video-coding framework for tangible and multi-touch systems. This framework allows researchers to rigorously categorize user actions as either pragmatic or epistemic, providing

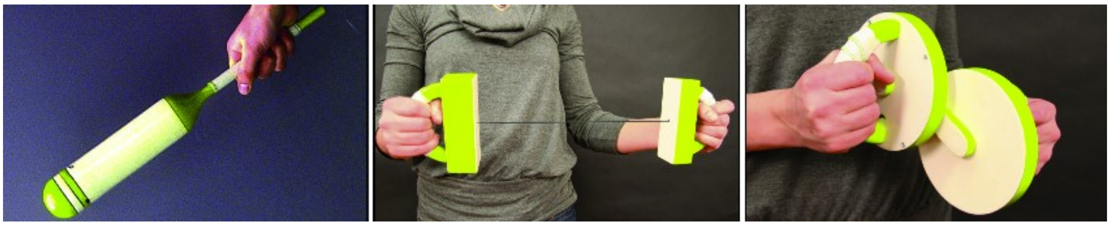

Figure 3.1. Three tangible artifacts from MoSo Tangibles [BHA10]. They exemplify how different *image schemas* can be used to teach children the sound concept of *pitch*. From left to right: the stick (low-high), the puller (near-far) and the rotator (slow-fast).

a glimpse of the deeper benefits of tangible systems. While groundbreaking, this framework still constrains epistemic actions to a single categorical label. As such, this dissertation will address the evaluation limitation of tangible interaction by further expanding on Kirsh and Antle et al.'s ideas: first, by developing our understanding of epistemic activity; and second, by applying that knowledge to novel evaluation tools for HCI.

TIREC: RECORDING TANGIBLE INTERACTION

4

As mentioned earlier, proponents of the tangible paradigm typically ascribe it a wide range of benefit properties over more traditional desktop/graphical interaction paradigms (e.g., improved support for collaboration). While this list is compelling and tacit support for the validity of these claims forms much of the basis for the field's growth, it remains highly challenging to perform the practical work of designing and developing effective tangible systems [SJ09]. How can beneficial qualities, such as supporting collaboration or learning, be targeted or maximized? How can alternative design candidates be meaningfully compared or specific claims about system performance empirically demonstrated? This chapter highlights a lack of evaluation methods specific to tangible interaction [SH10] and argues that this limitation substantially contributes to the absence of firm empirical evidence demonstrating the advantages of tangible interaction [A07, FTJ08b, N10] or concrete design guidance as to how to create tangible systems. Indeed, this chapter argues that continued development of the field tangible interaction will require the development of evaluation methods that are appropriate to the intrinsic qualities and challenges of the tangible paradigm – only by developing such tools will the benefits of tangible systems be fully understood and established.

As a first step to address this issue, this chapter introduces **tiREC** (Tangible Interaction Recorder), a novel tool that provides researchers and designers with an open, coherent, and practical way to systematically log and record events and actions in both tangible and desktop interactive systems. tiREC was created to enable systematic comparisons between designs candidates of particular tangible systems, between different tangible systems, and between tangible systems and other interaction paradigms. It also aims to facilitate the sharing and interpretation of empirical data between researchers in the field. This chapter reviews related work to establish a need for more systematic tools to evaluate tangible systems, introduces the structure of tiREC and describes a case study in which it was used to capture and analyze data: an interactive musical tabletop system. This chapter also describes how tiREC can be embedded in the frameworks and modelling languages specific to tangible interaction. Finally, the work presented in this chapter arguably represents a first step towards generating a specific evaluation method for the tangible interaction.

Evaluating Tangible Interaction

Despite existing a wide range of applications domains, from educational games [e.g., MKM07] to problem-solving tasks [e.g., UIJ05], it remains challenging to design and develop effective tangible systems [SH10]. This chapter argues that this is in part due to the fact that the in the past the majority of work in tangible interaction focused on developing proof-of-concept prototype systems [SJ09], with only a minority of articles providing clear documentation and design rationale [e.g., FT06]. Such rationale needs be critical, detailed and nuanced, covering broad aspects of a design from the form (e.g. size, shape, or colour) of the physical tokens used, to the fundamental mechanisms that underlie the interaction [JGH+08]. While these questions are likely considered during the design process of every tangible

system, this chapter argues these details are typically not reported and that this issue has been the focus of little dedicated and systematic research [H12].

tiREC Setup	
Method	**Parameters**
recordingMode	String rtype (*STANDARD, TUIML*)
addObject	String *oname*, *ttype* (*TOKEN, CONTAINER, TOOL*)
addConstraint	String *cname*, *ctype* (*SURFACE_INTERACTIVE, SURFACE_ON/OFF, OBJECT*)
addCustomEvent	String *ce1_name*
addCustomEvent	String *ce2_name*, String[] *vars*
Main Events Tangible, Touch and Mouse	
pickUp	String *oname*, *cname*, Int *x*, *y*
pickUp	String *oname*, *cname*
dropDown	String *oname*, *cname*
slideStart	String *oname*, *cname*, Int *x*, *y*
slideStop	String *oname*, *cname*, Int *x*, *y*
group	String *oname*, *cname*, Int *gid*
ungroup	String *oname*, *cname*, Int *gid*
customEvent	String *ce1_name*
customEvent	String *ce2_name*, String[] *vals*
Touch and Mouse-specific Events	
tap	String *cname*, Int *x*, *y*
pinch	String *cname*, Int *x*, *y*
click	String *cname*, Int *x*, *y*
pressStart	String *cname*, Int *x*, *y*
pressStop	String *cname*, Int *x*, *y*
TUIML Setup and Events	
addTAC	Int *tid*, String *var*, *token*
addObjectTAC	String *oname*, Int *tid*
addConstraintTAC	String *cname*, Int *tid*
addAction	String *aname*, *response*, *event*
addActionTAC	String *aname*, Int *tid*
pickUp	Int *tac id*, String *oname*, *cname*

Table 4.1. Example list of tiREC events. The comprehensive Javadoc can be found in *mysecondplace.org/tiREC.zip*.

This lack of focused, accrued design knowledge is aggravated by the lack of specific evaluation methods designed for tangible interaction. Typically, most evaluations of tangible systems adapt traditional methods from HCI [SH10]. These include a limited number of comparative studies, where researchers try to quantify the benefits of a tangible design by comparing it to other tangible designs and interfaces using metrics such as task completion times or error rates [e.g., CRR08] – which may not be able to capture the true richness or benefits of a tangible interaction. An alternative approach is to use observational protocols and video analysis, an approach that captures the full scope of behaviour with tangible systems, but is labour intensive and can be prone to novelty effects [ZAM+12], or inadequate for measuring long-term effects such as learning [SH10]. This chapter argues that evaluation techniques specifically intended to assess tangible interaction will help the field develop by enabling more effective comparisons between different tangible systems. A body of such evaluations will contribute to design knowledge in the field by allowing researchers to relate particular design features with particular user behaviours or performance outcomes.

Contrasting Interaction Paradigms

Comparative studies are the cornerstone of HCI. This chapter argues that comparative studies between tangible interaction and other paradigms are essential to isolate the real benefits and impact of tangibility. However, while such efforts are fundamental, they are highly challenging to design and conduct. One of the reasons for this is that it is hard to isolate the intrinsic characteristics of tangible systems that are important candidates for study. For example, natural support for collaboration is frequently cited as one of the hallmark benefits of tangible systems [e.g., ZAM+12]. The qualities that make tangible systems ideal for collaborative tasks are reported to include the ability to provide users with multiple access points which reduce interaction

bottlenecks and invite participation [HSC+09] and, to enable highly visible actions and parallel input among collaborators [HB06]. However, these qualities are not unique to tangible interaction and can also be found in other novel paradigms such as on tabletop multi-touch computers [WIH+08]. Moreover, systems that employ multi-touch input can also be expected to enable spatial multiplexing and bi-manual input – benefits also traditionally credited to tangible systems [FB97, KST+09] [10, 18].

Additionally, existing comparative studies that contrast tangible interaction with other paradigms are also affected by the methodological criticisms reviewed previously: they use simplistic metrics such as performance speed and efficiency [e.g., DPM+13], or they are 'unfair' in the comparisons they undertake by, for example, providing features in a tangible system that are absent in the tested graphical counterpart [e.g., SKM09]. These limitations make it currently challenging to both fully understand the design tradeoffs between different interaction paradigms, and to empirically demonstrate the benefits of developing computer systems using tangible interfaces. This chapter argues that these issues partly stem from the lack of field-specific tools to measure, compare and share user performance across systems employing different interaction paradigms.

Design and Development

tiREC is a programmatic framework developed with the goal of facilitating the gathering, interpretation and sharing of empirical data from comparative studies in the field of tangible interaction. In order to achieve this, several objectives were defined: tiREC had to be easy to integrate into applications, requiring a minimum amount of learning and little additional code; it needed to support custom events and work in different system setups (e.g. interactive tabletops, Arduino); it had to be easily integrated with current frameworks

and modeling tools; and, the data generated through tiREC had to be structured, concise and readable. In sum, tiREC aims to enable researchers and designers to coherently and systematically compare: different tangible design candidates; different study setups of tangible systems; and, computer systems employing different interaction paradigms. This chapter argues that the data gathered through tiREC will contribute to the creation of meta-knowledge regarding the inherent benefits and properties of tangible interaction, facilitating future work on the design of tangible systems.

tiREC was developed as a library for Processing, a Java-based, cross-platform, open-source programming language that is a popular option in most bespoke interactive tabletops (ITS), and available in several commercial ITS solutions (e.g., the ReacTable [JGA+07], Microsoft PixelSense). The information obtained through tiREC is recorded in a simple, self-contained, transactional SQLite database. This database is created with the aid of SQLJet, an open-source Java API for SQLite. The entity–relationship diagram can be visualized in Annex 1. The tiREC tool can be downloaded from *mysecondplace.org/tiREC.zip*. The user data logged with tiREC is structured to follow Holmquist et al.'s [HRL99] classification of tangible objects (into containers, tokens, and tools) and Ulmer et al.'s [UIJ05] classification of types of tangible interaction (interactive surfaces, constructive assemblies, and tokens+constraints). This theoretically grounded structure provides a framework for representing each physical object or token in a tangible system. Furthermore, each constraint can be modelled as either a token present/absent constraint, a spatial constraint (where X and Y position information extends basic presence information, as on an ITS) and a constraint created by one or more tangible objects (as in the Urp [UI99], where a physical building is both a token, and a constraint to other physical buildings). tiREC can also record time-logged events that occur when tangible objects are manipulated and when other UI elements are used. These events

include things as diverse as mouse clicks, taps and gestures on a surface, picking up and grouping of tangible objects and user-defined custom events. In order to keep tiREC platform independent, we note that it is the responsibility of application developers to identify such events and link them to the tiREC data logging system. For a summary of the events tiREC can record see Table 4.1 – the comprehensive Javadoc can be viewed in *mysecondplace.org/tiREC.zip*.

TUIML Integration

In order to facilitate the use of tiREC in the study tangible systems, it is also integrated with Shaer et al.'s [SJ09] Tangible User Interface Modelling Language (TUIML), a comprehensive and detailed modelling tool for specifying and defining the behaviour of tangible systems. TUIML provides explanatory power and system documentation through notation tools that emphasize the abstract, logical structure of the problem domain, design solution or application. TUIML's notation comes in the form of a high-level User Interface Description Language (UIDL) [O92] that describes a visual specification technique to represent tangible interaction as relationships between tokens (physical objects that represent digital content) and constraints (physical objects that constrain the behaviour of tokens) [UIJ05].

TAC	Representation			Association	Manipulation	
	Variable	Token	Constraint	TAC Graphics	Action	Response
1	Building	Building model	Surface, Other buildings	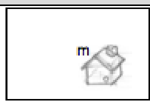	Add	Displays shadow according to time
					Remove	Removes related info from display
					Move	Updates display
2	Distance	Distance tool	Two buildings, Surface	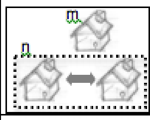	Add	Displays distance
					Remove	Hides distance
3	Wind Simulation	Wind tool	Buildings, Surface	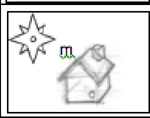	Add	Displays wind
					Remove	Hides wind
					Move	Updates wind

Figure 4.1. The TAC palette for Urp, a well-known tangible system for urban planning simulations [UI99].

These relationships are described as TACs and are represented in a TAC palette, a table that describes the grammar of all possible TACs in a tangible system design. Figure 4.1 presents an example of the TAC palette for Urp [UI99], one of earliest tabletop tangible systems. Table 4.1 exemplifies how tiREC can natively support modelling tools such as TUIML. By embedding the data it generates in these processes, it is the ultimate goal of tiREC to provide researchers with a structured design and evaluation tool based on objective data from empirical studies.

Case Study: jamTable

The case study presented in this chapter is a tangible musical sequencer based on a bespoke ITS. Named jamTable, it enables one or more musicians to interact with casual users controlling a tangible sequencer. The system's interface comprises of 24 music tiles (divided into seven columns), seven control tiles (one below each column), a record tile and two sets of physical objects: music and control tokens – see Figure 4.2. Interaction is as simple as placing tokens on tiles. Music tokens can either record an instrument's output in real-time when placed on the record tile or play back recorded sounds when placed on any of the music tiles. On the other hand, each control token can change the volume, the pitch or apply the popular drive effect to any music tokens playing in parallel in the column directly above. jamTable was developed using the Processing programming language, the reacTIVision tracking software and the TUIO messaging protocol. Lastly, it was deployed on a 120x70 centimetres bespoke ITS using the Diffused Surface Illumination (DSI) method.

The goal of this work was to provide researchers with initial insights regarding the real-world applicability of tangible musical applications in learning and performance activities. This was achieved by contrasting the

performance of trained musicians and everyday users (with no musical training) as they interact with playing musicians through the jamTable. The study itself followed a between subjects design based on these two setup conditions: two musician-musician (MM), and two novice-musician (NM) pairs. Study sessions started with the assignment of these roles by the researcher, followed by a small introduction to the jamTable and how it enables interaction with a musical instrument in real-time. After a five minute practice session, each pair of participants had 15 minutes to accomplish one single task: produce a music sequence to their satisfaction. Because this research was motivated by an open and exploratory research question – the suitability of tangible interaction to support musical learning and performance – it was necessary to gather both qualitative and quantitative measures. Quantitative measures were especially important to help gauge the differences in the performance and experience of MM and NM pairs. As such, tiREC was deployed primarily to facilitate the retrieval of consistent interaction data between the two conditions. This data included: the highest number of music tokens used simultaneously; the number of control tokens applied; how many unique control tokens were used; the amount of time these control tokens were active; the number of recordings performed; and the duration of such recordings.

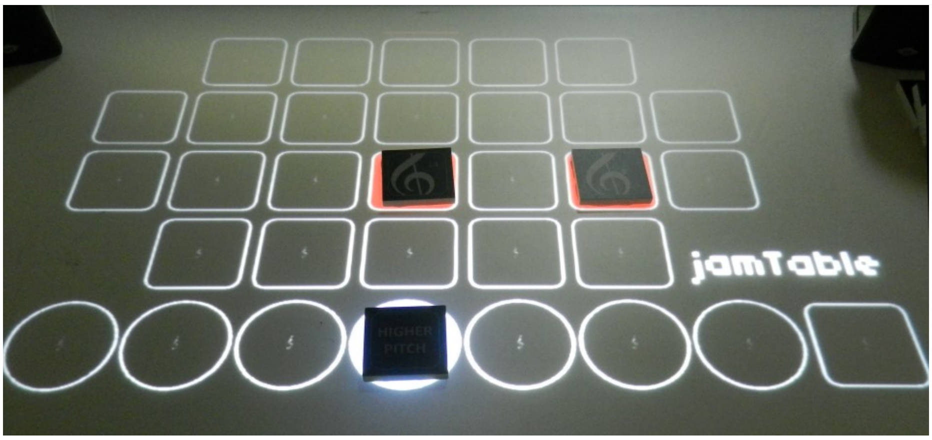

Figure 4.2. The jamTable: a tangible musical sequencer comprising of 24 music tiles, seven control tiles below and a record tile on the bottom right corner. In this figure two music tokens are playing in sequence (highlighted in red), and a control token is raising the pitch of the leftmost music token (highlighted in blue).

Several different options were considered when deploying tiREC with the jamTable. Firstly, the system deployed over 60 physical objects: 30 music tokens and 35 control tokens. Due to the high number of tokens available for recording, each was identified in tiREC with their corresponding TUIO ID instead of individual, personalized names – e.g., *addObject("musictoken_125", "TOKEN")*. Secondly, and even though jamTable was implemented on an ITS, several present/absent constraints were defined in place of a single unique spatial constraint representing the entire surface – e.g., *addConstraint("musictile_1x2", "SURFACE_ON/OFF")*. Each of these constraints represented one of the 32 tiles present in the interface (music, control and record tiles). Once this setup was in place, tiREC simply recorded the picking up and dropping down of specific tokens within specific constraints. This made it particularly easy, and less error-prone, to examine and interpret the recorded events and interaction. The source code describing this tiREC implementation can be viewed in *mysecondplace.org/jamTable.zip*.

The data gathered through tiREC was particularly useful in the generation of initial insights regarding the suitability of tangible systems for musical teaching in collaborative scenarios. This data demonstrated, for example, how NM pairs were largely more open to explore the effects of control tokens and how the jamTable enabled participants to learn different musical concepts – from tracks to arrangements to riffs. While this data was studied in tandem with more qualitative measures, the key empirical data in this research was captured through tiREC – a step forward in the empirical validation of several of the benefits of tangible musical applications. Additionally, tiREC generated empirical data attesting to the unique benefits of tangible interaction for collaboration: unlike touch-based ITS that equally provide visibility of actions and system state when all users are directly interacting with the system, tangible applications can provide the same level of support even when one or more users do *not* have direct access to the system (such as the playing

musician interacting with the jamTable). Again, this information is broadly applicable to any researcher or developer designing tangible systems for two of the most popular application areas in the field – public installations and musical performances. Finally, tiREC is also a suitable tool for generating empirical data to help contextualize qualitative data gathered from different sources. Several of the insights discussed in the jamTable research were gathered through a combination of tiREC data with data from questionnaires and video recordings from the users' performances.

Future Work and Conclusions

This chapter argues that the field of tangible interaction needs to develop and adopt evaluation methods that reflect its unique characteristics. As such, one of the goals of this dissertation is to position tiREC as tool that can provide a first step to the creation of empirical evaluation frameworks that are specific to the field of tangible interaction. Prior reviews of the field suggest this effort is novel. For example, Mazalek et al.'s [MH09] comprehensive summary of tangible frameworks tackles issues such as abstracting, designing and building tangible systems. While such tools provide conceptual guidance and descriptive power, they are far from being embedded in the iterative development cycle of tangible systems. tiREC addresses this lack by natively supporting existing frameworks, such as the TUIML [SJ09]. However, much more work is required to develop a mature evaluation framework specific to tangible interaction. This includes defining methods and tools the can further integrate the tiREC with modelling languages such as TUIML. This chapter identifies the development of automatic systems for generating and annotating a system's TAC palette with usage data recorded using tiREC as a fruitful area for future work. Additionally, in order to create a fully open, generalizable and usable system, several technical challenges need to be addressed in future work. These include:

- Opening tiREC development to the community, so as to encourage the broadest deployment, customization and generation of tiREC data.

- Developing versions of tiREC in additional programming languages such as C# or Arduino, so as encourage adoption and facilitate deployment.

- Integrating additional input hardware such as the Kinect, the Leap Motion, or simple voice recording. Adding such functionality will enable the generation of broader sets of user events, such as where users position themselves around a ITS, what events take place directly above the surface of an interactive tabletop (e.g., the passing of physical objects between users), or common bodily actions such as verbal shadowing and gesturing.

- Creating an online repository open to the research community for the sharing of study designs, captured tiREC data, and custom visualizations created.

This chapter argues that accomplishing these goals will pave the way towards a systematic, rigorous, and much needed tool for empirically evaluating and sharing data about tangible interaction. Both the tangible systems presented in chapters 5 and 6 were studied with the help of tiREC.

EPISTEMIC ACTIONS IN ACTION

5 This dissertation introduced tangible interaction as interfaces that facilitate what is known as *tangible thinking* [SH10]. While several authors have argued this concept is poorly supported by purely virtual systems [e.g., SKM09], there is still a lack of firm empirical evidence demonstrating the concrete advantages of using tangible interaction to help users deal computer mediated tasks [A07, FGH+00, M07, G91, FTJ08]. This dissertation argues that, in order to effectively challenge traditional interface paradigms, an improved understanding of the benefits of tangible systems needs be established. An important aspect of achieving this is rigorous, repeatable and equivalent comparisons between interaction paradigms.

Moving towards this objective, this chapter presents a comparative study between three different interfaces to the Four-in-a-row board game (see Figure 5.1) and explores how different interfaces to this task affect how users perform it. In this game, players take turns dropping coloured disks in a vertical grid, with the goal of connecting four disks of the same colour either horizontally, vertically or diagonally. Two of the interfaces studied in this work show both the grid and the disks in a graphical display, differing only in the input method: mouse versus direct touch. A third interface is composed of an augmented physical game-board and real disks. These three systems were careful designed to ensure interactions were functionally equivalent;

and a user study is presented that looks at both quantitative and qualitative data on user performance. The results highlight the relevance of projection and **epistemic actions** to this task and suggest that the different interface forms afford instantiation of these activities in different ways.

Method

Experimental Design and Participants: The study presented in this chapter follows a within-subjects repeated measures design based on three interface conditions: tangible, touch and mouse. In total, there were 36 participants: 22 males and 14 females. 19 participants were from Europe, 14 from Asia, two from North America and one from South America. Their ages ranged from 16 to 34 (M = 24, SD = 4.16), and with the exception of one, all participants were students at local universities. Of the 36 participants, only four had never played Four-in-a-row before.

The participants completed the study in groups of three (for a total of 12 sessions) and also completed a total of three game sessions, one using each of the interface conditions. To mitigate potential practice or fatigue effects, the order in which the conditions were experienced was fully balanced – two groups completed each of the six possible order conditions. All participants received compensation in the form of a 5€ voucher valid at a range of local stores and service providers. Success at the game was also rewarded – the participant who won the most games in each group received an additional 5€ voucher while the top three participants in the whole study received a further 10€ voucher.

Figure 5.1. A common game of Four-in-a-row [B13].

Materials: Four-in-a-row is a board game where two players take turns dropping coloured disks in a 7x6 vertical grid (see Figure 5.1). The objective is to be the first player to connect four disks of the same colour in either a horizontal, vertical or diagonal line. Disks are dropped into columns from the top, meaning that the gradually changing accumulation of the disks in the different columns is an important game play element. For the purposes of this study the Four-in-a-row game mechanics were altered. Essentially, a third player (basically a third piece colour) was added, ensuring more complex changes took place between each player's moves (the introduction of two rather than one new piece). This also ensured that no users had specific prior exposure to the game dynamics, as they were as least partially novel.

Interactive feedback was also introduced to the game system. This took the form of highlighting in response to exploratory gestures with the game pieces. Essentially, if participants positioned a game token at the top of one of the game board columns for a dwell period in excess of one second they were presented with appropriately coloured visual feedback indicating the position the disk would reach if dropped (see Figure 5.2). This feature was termed *hovering* feedback, and it was intended to provide information on the board's possible future states prior to making an actual move in the game – in essence an epistemic action.

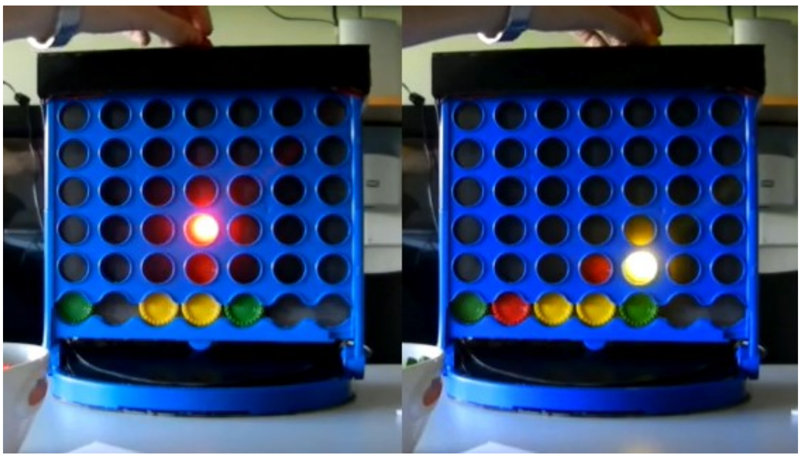

Figure 5.2. The tangible version of the game, showing the hovering feature for both a red and yellow disks (from left to right).

Three versions of the game were produced to support tangible, touch and mouse interaction. Each featured a 7x6 grid of holes with a total visible size of 26x24 cm. Each used game disks of 3 cm in diameter (0.5 cm thick for the tangible version) that could be moved to locations directly above the game board's columns to receive the hovering feedback and/or to be dropped into place. The mouse and touch versions used fully graphical interfaces developed using the Processing programming language and displayed on a small portion of a 120 cm vertical flat screen. Mouse input was provided via a standard peripheral attached to the computer driving this display, while touch input was achieved via a SMART Board Interactive Display Overlay placed in front of the screen. In both these interfaces, simply clicking the mouse or touching the screen caused a drag-able icon of a game disk to appear under the cursor (or finger). This could be positioned directly above the board to gain access to the hover feedback or released there to add a piece to the game. The tangible version, on the other hand, was based on a physical game board into which physical disks were placed. The hover feedback was realized via two vertically stacked photo interrupters mounted on top of each of the columns (14 sensors in total). Placing a physical token in between the top emitter and sensor triggered the hovering feature (see Figure 5.2), while

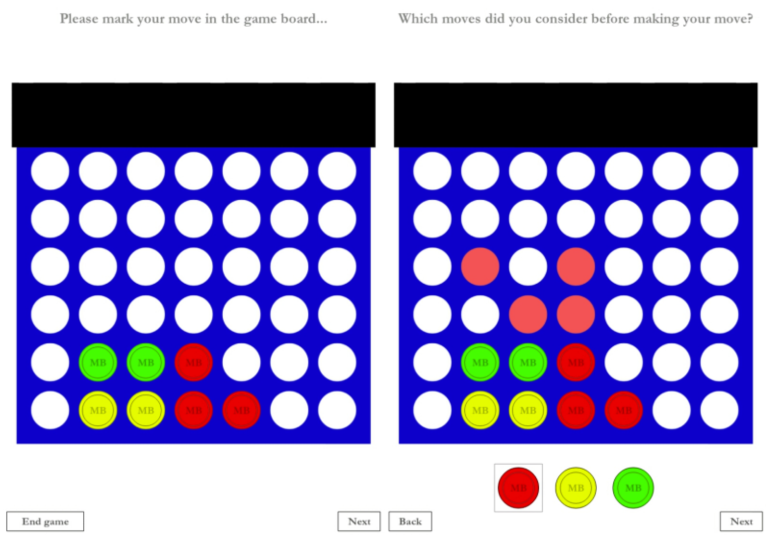

Figure 5.3. The Four-in-a-row tablet application, which helps participants represent their mental projections after each play. On the left, the board is updated with a red disk. On the right, the reasoning is given for playing the red disk on that location.

an interruption of the bottom sensor was used to indicate a disk drop. Each of the bottom sensors was located 0.5 cm above the game board, with the top sensors located at 1.5 cm. Graphical feedback for the hover event was enabled by placing a diffuser screen (Rosco Grey) and seven strips of digitally addressable RGB LEDs behind the board (so that there was one LED for game-board hole). All electronics were connected to an Arduino Mega microprocessor that monitored input and displayed the feedback. This construction ensured a bright, responsive display and that participants were only able to see the board from one side (as in the case of the two other versions of the system).

Procedure: In each session of the study a group of three participants played three games of Four-in-a-row against each other, one game in each of the three interfaces. Sessions commenced with brief introductions among the group, an explanation of the condition sequence, game rules and compensation structure, and assignment of each participant to a disk colour for the duration of the study (red, yellow or green).

The experimental interfaces were all presented in the same small and otherwise empty office. Each of the three games followed an identical structure: the three participants were invited to interact informally with the interface (maximum of five minutes) placing disks and becoming acquainted with the hovering feature. They were then asked to move to an adjacent room, where chairs and snacks were provided. Whilst there, they were instructed not talk about the game. Participants entered the game room individually in order to make their moves, ensuring that their epistemic actions were private. To decrease study and system complexity (colour sensors), the first player to move was always randomly selected but the sequence of players was always the same: red, followed by yellow, followed by green. Information reminding players of this sequence was prominently displayed in both game and waiting

rooms. Participants completed a range of subjective measures between individual turns and at the end of each game, as detailed next.

Measures: In addition to general game play results, the metrics used in the study included:

- *Time to play* – As used in similar problems [e.g. K09b] this metric is defined as the total amount of time a participant takes to complete his or her turn. The start point of this period was calculated by equipping all three versions of the game with a face recognition system consisting of a standard webcam and the Processing OpenCV computer vision library. When a participant faced the game board, this event was recognized, a sound played and the initial time logged. The period ended when the participant dropped a disk into the grid, as detected by the game software. In order to ensure that the time taken to pick up or select a disk did not influence this measurement, participants in the tangible version started their turn with a disk already in hand. Similarly, participants in the two graphical versions "picked up" a disc simply by clicking the mouse or touching the screen, irrespective of where these events occurred. Finally, this metric was recorded with tiREC (chapter four).

- *Epistemic actions* – The number of the epistemic actions performed in each of the three interfaces was recorded in two ways: automatically, by recording when the hovering interface feature was triggered (tiREC) and; through video analysis (two observers, with a high inter-rater reliability – a Kappa of 0.701 [LK77]). In this latter case, epistemic actions took the form of pointing gestures at or in front of the game board. These were divided into those made with or without the game disk (e.g., Figure 5.5).

- *Mental projection* - In the third and final game of each session, and directly after completing each turn, participants used a tablet application to explain and justify their moves. This was achieved via a custom Android app that showed a Four-in-a-row game board and enabled them to tap grid cells to illustrate not only the current state of the game, but also the potential moves they considered whilst planning their play (see Figure 5.3). This application was deployed on 10.1″ Android tablet and was developed using the Processing programming language. The application logged two key data points: both the number of candidate positions participants considered for their move and the number of possible opponent responses.

- *Subjective Workload* - Each participant completed the NASA TLX, Hart and Staveland's six-item workload questionnaire [HS88] at the end of each game.

Results

The experimental results are presented below. Unless otherwise noted, all analyses were conducted as repeated measures one-way ANOVAs over the three experimental conditions (tangible, touch and mouse). Greenhouse-Geisser corrections were used when assumptions of sphericity were violated and all *post-hoc* comparisons were t-tests with Bonferroni corrections.

Time to Play: The mean time to play across all interfaces is presented in Table 5.1. Outliers resulting from problems (e.g., jammed disks, the system failing to detect a new turn) with the tangible version of the game were removed prior to analysis. A significant trend was found in this data (F (2, 52) = 8.202, p = 0.001) and subsequent pair-wise differences were revealed between the tangible and mouse interfaces (p = 0.004), but not between tangible and touch

(p = 0.160) nor, although there was an observable trend, the mouse and touch conditions (p = 0.067).

Epistemic Actions: The mean results for epistemic actions performed are presented in Table 5.1. These are divided in the actions that relied on the hovering feature and those that relied to pointing with and without a disk in hand. Data from the hovering feature was only considered when the feedback light was on for at least one second. Significant trends were observed for the use of the hovering feature (F (2, 52) = 8.772, p = 0.001) and the gestures without disks (F (1.220, 20.748) = 20.061, p < 0.001) but not for gestures with disks (F (1.173, 19.936) = 1.740, p = 0.204). For both of these two significant trends pair-wise comparisons revealed differences between tangible and touch and touch and mouse conditions (p < 0.005 for both). No differences were found between the tangible and mouse conditions (p = 0.489). Finally, all participants were observed performing epistemic actions. Across a full session the minimum number of such actions performed by a participant was three, the maximum 41.

	Tangible	Touch	Mouse
Time to play	17.85 (5.81)	22.89 (12.27)	27.32 (13.73)
Own	6.25 (4.35)	7.50 (3.78)	13.08 (7.95)
Opponents	3.08 (3.48)	1 (0.63)	3.08 (4.12)
Hovering feature	5.37 (4.46)	2.41 (2.45)	5.56 (4.64)
With disk	1.11 (1.68)	0.56 (0.78)	1.56 (2.15)
Without disk	0.50 (1.20)	2.22 (2.05)	0.28 (0.75)
Epistemic Actions p/ game	6.98	5.19	7.4
Total Drawn Games	7	2	8

Table 5.1. **Time to play:** mean time to play in seconds according to each game interface (12 games per interface). **Mental projection (own / opponents):** self-reported mean number of moves considered prior to play (four games per interface – the tablet application was only used in the last game of each session). **Epistemic actions (hovering / pointing):** mean occurrence rates for epistemic actions recorded with the hovering feature and obtained through video analysis. **Epistemic actions vs drawn games:** Mean occurrence rates of epistemic actions per game vs number of games drawn for the three interfaces. Numbers in brackets are standard deviations.

Mental Projection: The mean results from the data recorded with the tablet application are presented in Table 5.1. Data was only considered after each player had made two moves, to ensure some degree of game complexity. The number of own moves considered varied significantly (independent samples ANOVA, F (2, 27) = 4.19, p = 0.026) and pair-wise comparisons showed the significant changes to be between tangible and mouse (F (11, 11) = 3.34, p = 0.029), and touch and mouse interfaces (F (11, 5) = 4.86, p = 0.047) but not between tangible and touch (p = 0.355). There were no statistically significant differences in the number of opponents' moves recorded (p = 0.430).

Subjective Workload: The mean results from the TLX workload questionnaire are presented in Figure 5.4. Analyses were conducted on overall workload and each individual scale. No significant differences were observed (p > 0.177 in all cases).

Gameplay Statistics: The mean results for number of disks used and number of victories per interface is presented in Figure 5.4. These show that in average, 26 (SD = 10.82) disks were used per game in the tangible, 25.2 (SD = 9.22) in the touch and 30.33 (SD = 7.57) in the mouse version of Four-in-a-row. A full board contains 42 disks. Additionally, five games ended in a victory in the tangible version (seven draws), 10 in the touch version (two draws), and

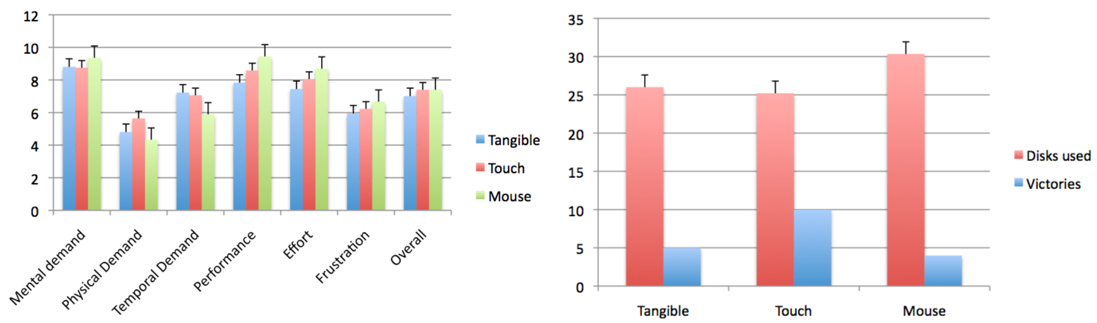

Figure 5.4. Left: TLX data per condition (0-20 range). Lower scores represent lower perceived workload. Right: general game results. Mean disks used in red (of a maximum of 42 per game) and number of games ending in a victory in blue (of a total of 12 games per interface). Bars show standard deviations.

four in the mouse version (eight draws). No participants won all three games, and only on one occasion did a single participant win twice in a session. Finally, 47.37% of the participants made more epistemic actions than their opponents when winning a game, and 43.75% took more time to play.

Discussion

The study presented in this chapter is grounded on Kirsh's work on mental projection [K09b], which investigated the time that players took to make moves in three different representations of the game of Tic-Tac-Toe. By showing that users played faster when game materials and elements were visible (such as the grid board or the X and Os on that board), Kirsh's work demonstrated the importance of having *out-of-the-mind* structures on which to anchor cognition. Moving beyond these findings, this chapter's goal was to study if a physical, tangible representation of a task could better serve as anchor for users' cognitive endeavours compared to (as otherwise similar as possible) graphical counterparts.

The basic experimental data are ambiguous on the benefits of the tangible interaction style over two virtual versions. The time to play data showed one significant difference – the tangible interface improved over mouse input but not touch-screen. However, this can be partly explained by prior authors' assertions that dragging objects with a mouse is more time consuming than

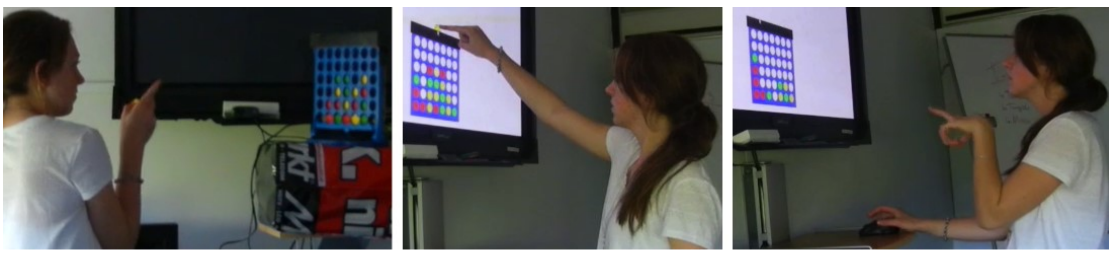

Figure 5.5. A participant playing Four-in-a-row in each of the interfaces (tangible, touch and mouse, respectively). In the tangible and mouse interfaces the participant is performing a pointing action without a disk, while in the touch interface he is triggering the hovering feature.

performing the same action through more direct input methods available in tangible and touch interaction styles [A12]. Furthermore, while there are noticeable differences in the data relating to mental projection of other's moves, it is worth noting that this comparison was subject to the influence of individual differences (i.e. it was not within-subjects). Furthermore, an ANOVA on the overall summed number of mental projections did not yield a significant result ($p = 0.068$). An alternative explanation for the increased rates of the reported projections in the mouse interface is also simply that the increased time it took to complete a move gave participants more time to spend thinking about the game-board.

These results cast doubt on the value of physical representations in computer-mediated tasks. However, before dismissing them, it is worth extending this discussion to include a more in-depth consideration of the cognitive work involved in such activities. As discussed previously, epistemic actions allow users to reduce the cognitive cost of maintaining or extending mental projections, as they allow users to externalize aspects of these artefacts – to make parts of them real [K09b, KM94]. This is a key factor underpinning the notion of thinking with things, a process that Kirsch characterizes as: "knowing what you are thinking by seeing what you are saying" [K09b]. A modern example of the importance of epistemic actions in problem-solving tasks comes from the Chess game available for the Microsoft PixelSense tabletop computer. An early version of this game only permitted valid moves but, in response to user feedback, an update was released that allowed for unconstrained and exploratory moves like those that can be performed on a normal chessboard.

Reflecting the importance of epistemic actions, this chapter focused on three possible behaviours: the hovering feature, as it enabled an explicit and automatically recordable epistemic action and pointing or touching the game

board with and without a disk, a general epistemic technique that is reported to helps users focus attention through symbolic marking [K95b]. Variations in the occurrence rates of these three actions in the study were complex. In the case of disk-hovers, rates were down in the touch interface while for unencumbered pointing, they were up. Conversely, no differences were observed for gestures when a user was holding a disk. These data stand in contrast to claims, typically based on the ease with which physical tokens can be grasped and manipulated, that tangible interaction is more suited to supporting epistemic actions than graphical interfaces [e.g. PI00]. Rather it highlights that epistemic actions are readily achievable in the digital domain, a suggestion supported by Kirsh's seminal work introducing the concept using a traditional and purely virtual version of the Tetris video game [KST+09]. Indeed, other authors have remarked on the diverse and flexible mechanisms by which people achieve and employ epistemic actions [e.g. KM94, WVV11].

This perspective helps explain the fact that while some of the epistemic actions were more commonly performed in the tangible interface (e.g. twice the number of hovering actions were observed compared to the touch interface), others were equivalently or less frequently performed. This suggests that peoples' need to employ epistemic actions leads them to take advantage of whatever resources are optimal to achieve this end. Furthermore, it is clear that the different interfaces and representations of the problem afforded different actions – gesturing with a coin in hand was arguably simpler in the tangible interface than in the touch interface, where one would first have to come in contact with the touch screen before the gesture could be achieved.

Another key question of interest is whether mental projection and epistemic action positively influenced user performance across the study as a whole. In

order to fully consider this issue, it is worth noting this takes place against the novel aspects of the game play. Introducing a third player (and consequently an additional colour of disk) to the game, while maintaining the board size and winning criteria (a line of four in length), made it substantially more challenging. This can be seen in the relatively even distribution of wins throughout the study – normally occurring individual differences in skill levels at the original Four-in-a-row game had little influence on outcomes in the experimental game. Indeed, only one player managed to win more than one game. In light of this, a draw was interpreted as representing a balanced game, where participants successfully predicted and prevented their opponents' plans. Although too interpretative to be subjected to formal statistical analysis, the number of draws did vary with the number of epistemic actions as shown in Table 5.1. This tentative relationship suggests that epistemic actions played a valuable role in helping participants understand the state of the game.

This idea is further supported by an analysis of the activities of winning participants. Essentially in 47.37% of wins, the victor was the participant who performed the greater number of epistemic actions. Considering the chance rate of adopting this position is equal among players, or 33% in the game studied here, this fact suggests users who were performing more epistemic actions were more likely to win regardless of interface they were using. This assertion provides tacit support for the tangible interface – although the total number of epistemic actions did not vary from interface to interface, the time data suggests they were performed substantial faster in the tangible game – 6.98 actions in just 17.85 seconds versus 7.4 in 27.32 seconds while operating the mouse driven interface. This result suggests that a key advantage of tangible interfaces may be that they have the potential to support rapid execution of useful and informative epistemic actions.

To conclude, meaningfully comparing the influence of a physical interface on a computer-mediated task is challenging. Tangible interaction naturally lends itself to 'unfair' comparisons as it offers interface features that cannot be effectively matched to graphical counterparts [e.g. SKM09]. As such, the main concern when developing the three versions of Four-in-a-row for this study was to ensure they behaved and responded consistently and equivalently. A range of techniques were used to achieve this, including removing the impact of disk selection by having participants start their turns in the tangible interface with a disk already in hand and simply clicking the mouse or touching the screen to summon a drag-able disk under a users cursor or finger in the graphical conditions. Furthermore, participants were constrained to releasing the disks in the same set of viable positions at the top of the game board in all three interfaces. The study was also explicitly designed to promote and isolate a clear and observable projection-action-projection cycle by having participants leave the game room in-between turns. This ensured they faced an evolved board state afresh every time they needed to play a piece. This chapter argues that only with such epistemic measures in place is it possible to meaningfully compare performance and attempt to understand the effects of the physicality of a representation on people's mental efforts to solve the problem at hand.

The results of this study were informative initial steps towards this goal, but many questions remain unanswered. One of the most important relates to the fact that mental projection is not a free process [K09b]. In the context of the game studied here it requires the anchoring of *imagined* game disks to the *real* game board. Exploring the trade-off between this cost and the benefits it brings is a good topic for future work. Furthermore, it has been reported that the usefulness of anchors depends on both a person's visualization abilities and the overall complexity of the problem at hand [K09b]. Exploring how tangible representations influence this threshold, essentially one of the point

at which mental projection and epistemic actions become *profitable* would likely be fruitful and interesting.

In summary, the work presented in this chapter highlights the difficulty of formally demonstrating value in the tangible interaction paradigm. It equips itself with an appropriate toolbox composed of a theoretical proposition, a carefully designed set of alternative systems and a hypothesis for how these will interact. However, although the results hint at advantages of the tangible approach, few direct effects were observed. The study did successfully highlight the importance of the theoretical ideas on which it was based – mental projection and epistemic actions – and the results indicate that the form in which users will instantiate these concepts systematically varies from interface to interface. This suggests that future work on tangible systems should focus on understanding and designing projection techniques that are well matched to instantiation in the physical world. As this work goes forward, continued efforts to meaningfully compare tangible systems with other interface paradigms will be required to understand the advantages this paradigm has to offer in terms of support for epistemic actions in computer-mediated tasks.

OFFLINE EPISTEMIC ACTIONS

6

This chapter continues the exploration of how tangible systems can be understood and designed so that *tangible thinking* can effectively occur during interaction. This is a challenging problem – embeddedness in the world has numerous (negative and positive) implications. For example, recent work has highlighted the inevitability of manipulating tangible objects beyond the scope of a system's ability to sense them [e.g. FTJ08]. Negative implications of this behaviour include how it can frustrate and confuse users (e.g., when metaphors between object manipulation in the real and digital world break down [H12]). On the other hand, in tabletop displays how users organize and manage tangible objects outside of the sensing surface of the system has been reported to positively impact user behaviour and system use (e.g., by helping to seamlessly mediate supportive actions such as token passing in collaborative play [SSV+11]).

This chapter argues that such peripheral, non-task centric activities represent a practically important and potentially rich design space capable of supporting a range of high-level tasks. More importantly, as in real-world tasks, the majority of these activities can be understood as epistemic actions (e.g., how participants often pointed with a game disk while playing Four-in-a-row). Compelling examples illustrating this point exist. For instance, in

Durrell-Bishop's seminal Marble Answering Machine [A99] glass beads linked to voicemail messages could be placed in physically customizable passive containers associated with specific family members or activities (see Figure 6.1). Similarly, in Mementos [EO10], simply selecting a set of physical tokens to use before interaction was characterized as a collaborative decision making process; while the large number of tokens in the LogJam Video Logging system [CWP99] encouraged users to systematically structure and organize tokens away from the sensing surface. Taken together, these examples highlight the potential value and diversity of designing for manipulations of tangible elements that are not interpreted by computer systems, but also the piecemeal nature of research on this topic. This type of epistemic activity will be known for the rest of this dissertation as offline epistemic actions – actions that this dissertation defines as those that are *invisible* to computer systems.

This chapter suggests that offline activities represent an important, unique and unexplored aspect of tangible systems. It seeks to provide an early description of the impact of such activities through the development and evaluation of **Eco Planner**, a tangible, calendar-like tabletop application that requires the use and management of a large token set that represents routine activities such as doing the laundry or driving to work. Two studies of the

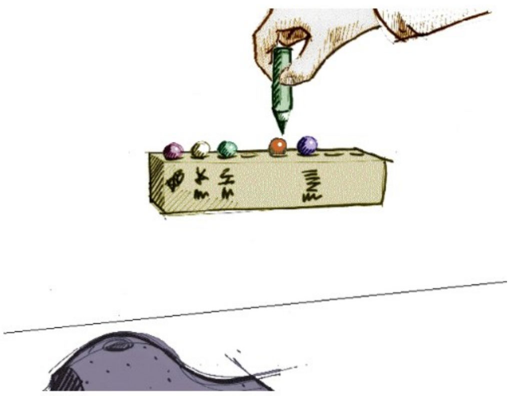

Figure 6.1. One of the possible offline activities in the Marble Answering Machine [A99]: tagging the physical marbles outside of the system according to different callers. In this example the box where the marbles are tagged and categorized is completely *invisible* to the system.

system are presented covering individual and paired system use (to highlight the offline use of physical tokens in collaborative action). The qualitative and quantitative results from these experiments are then synthesized into a set of design insights that focuses on the form, size and location areas where offline epistemic actions can occur. These guidelines are intended to frame future efforts in tangible interaction design so these systems better support what this dissertation argues is their core characteristic: their ability to effortlessly support epistemic actions that lower the cognitive demands of the task.

Eco Planner

Eco Planner was developed to explore the design of different areas that could better support offline epistemic actions. Four requirements guided the development: the need to deal with a broadly applicable and understandable problem domain – representing a set of activities on a grid or calendar-like schedule (as in SenseBoard [JIP+02]); the need to support a large token set, ensuring that users would have to manage tokens that could not always be in sensing range of the system; the need to provide users with different areas that could support offline activities; and the need to support collaboration, by allowing multiple tokens to be used simultaneously and for the overall task to be meaningful as a group activity.

Figure 6.2. A set of physical tokens from Eco Planner application. Each represents a household activity such as watching a movie, commuting to work or doing the laundry.

Interaction with Eco Planner is simple: users position iconic tokens on a grid to create routines representing their household activities. The focus of the interaction is on creating and visualizing activities that involve energy consumption in order to calculate and understand the environmental impact of

particular behaviour patterns and consumption choices. The interface is composed of a set of physical tokens that are set up on an interactive tabletop surface (ITS). Via iconic labels (see Figure 6.2), each token represents a specific activity (e.g., watching TV, doing the laundry), and users create their household routines by configuring these objects on Eco Planner's tabletop surface. The application provides real time feedback and recommendations relating to the objects arranged on its surface, including an overall score representing resources consumed by the entire schedule. Lastly, users can set the system to display these cues as either ecological or financial messages (e.g. by displaying the CO^2 footprint or expenses accumulated).

The Eco Planner interface is shown in Figure 6.3. The interface is divided in three major areas. The *routine* area takes the form of a large rectangular zone that represents a single day (from 7am to 11pm). Time is represented along the horizontal (or long) axis – tokens placed on the left of the routine area represent activities conducted in the morning, while tokens on the right represent activities in the evening. Tokens aligned vertically (or along the short axis) represent concurrent activities. The interface also features an *options* area, a small central zone where users can position tokens and interact with a simple interface that lets them commit to different behaviours relating

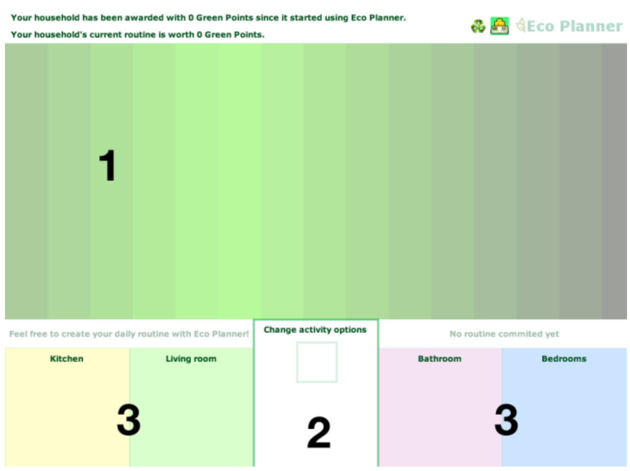

Figure 6.3. The Eco Planner interface: (1) the *routine* area, where users can spatially lay tokens to create their daily routines; (2) the *options* area, where users can commit to a series of changes to a particular activity; and (3) the *repository* area, where tokens can be rested.

to that activity – e.g., with the laundry token, users can choose to commit to always doing the laundry with a full load of clothes. To explore how spatial consistency is reflected in tangible interfaces, two versions of Eco Planner were developed with slightly different repository areas for tokens that were not in use: labelled and blank. In the labelled version (Figure 6.3), Eco Planner has four colour-coded repositories, each marked as a particular area of a home (kitchen, living room, bathroom and bedrooms). These labels provide users with a coherent and consistent location to store and rest unused tokens, e.g. they are able to stash tokens representing kitchen activities in the kitchen zone, thus supporting subsequent easy access for themselves or other users. In the blank version, the four repositories have a single colour (grey), providing users with no clues in how they should organize tokens outside of the system. In both cases, the repository zones were situated at the corners of the tabletop, sites where users have previously been observed storing unused tokens [SCI04].

Eco Planner was implemented using the Processing programming language and runs on a tabletop system developed using the Diffused Surface Illumination (DSI) method, the ReacTIVision tracking software and the TUIO messaging protocol. The interactive surface is 78 by 57 centimetres in size and composed of a diffusing layer (5mm thick Evonik ACRYLITE 7D006) placed on top of a sensing layer (10mm thick Evonik Endlighten). The sensing layer is wrapped with a string of 850nm IR diodes and positioned above a near throw projector and IR camera so that both can address the full area of the interactive surface – this DSI setup offers a simple and reliable way to track both fiducial markers and fingertips on an interactive tabletop. The tokens used in the system are 7 by 7 by 2.5 centimetres wooden cuboids with iconic labels affixed to their uppermost surface and ReacTIVision markers on their base (see Figure 6.2). 27 tokens were deployed in the Eco Planner system, representing 13 unique activities.

Method

This section describes a user study with the Eco Planner tabletop system. It follows a between groups design with three independent variables: collaboration (single/paired), to explore how unused tokens might support group activities; initial token configuration (local/distant), to explore how the initial configuration of tokens cues and influences the management and structuring of these by the users; and repository zone format (labelled/blank), to explore to which extent do users rely on such cues for spatial consistency when engaging in offline activities on tangible interfaces.

The collaboration and token configuration variables were paired such that single participants used distant tokens and paired participants used local tokens. Repository zone format was distributed equally across these two groups: half of the single/distant group and half of the paired/local group used the labelled repositories and half used the blank repositories. The collaboration variable was adjusted by recruiting two participant pools, individual users and pairs of users, as described in the participants' section of this chapter. Paired users were required to complete the experimental tasks collaboratively. The repository zone variable was modulated via enabling the two zone types in the Eco Planner interface as required. Finally, token configuration was varied (local/distant) via changes to the experimental procedure, as noted in the procedure section.

Participants:

- *Single users* – A total of 12 single participants were recruited, eight males and four females. There were five Portuguese, three Indians, two Venezuelans, an American and a Pakistani participant. Their ages ranged

from 23 to 35 (*M* = 26.5, *SD* = 4.2) and all were students or full time research staff at the local institute.

- *Paired users* – Six pairs of participants were recruited (four males, eight females), each living together for a minimum six months at the time of the study. There were four participants from the USA, two Taiwanese, two Portuguese, a Canadian, a Swede, a Kenyan and an Indian participant. Their ages ranged from 23 to 33 (*M* = 26.2, *SD* = 2.9), and all of them were either undergraduate or postgraduate students at the local institution. One pair of participants was a couple, whereas the other five were cohabitants (flat/roommates).

Procedure: The experiment started with participants being shown the tabletop display and receiving instructions on how to operate Eco Planner. All participants were then asked to model three different routines. The first two were derived from short written scenarios depicting two fictional households, while the final routine was based on participants' own activities at home and aimed at exploring differences in behaviour for more experienced users. Participants had total freedom over defining the routines. There was no time limit for the tasks and at the end of the experiment participants filled in a paper questionnaire.

The experimental procedures varied according to two types of initial token configuration. In the *local* condition, the physical tokens were initially set up in the repository areas of the Eco Planner system and were not returned to pre-set positions in-between experimental tasks. Participants were required to manage the tokens themselves as they created the three different routines. The goal of this setup was to explore how much users rely or are constrained by the initial spatial configuration these items, with special attention to these cues in a collaborative setting. On the other hand, in the *distant* condition an

experimenter set up the tokens on a desk a few meters from the Eco Planner tabletop. Participants were required to choose a set of tokens from the desk and bring them to the tabletop prior to interaction. After each routine was completed the tokens were returned to the distant desk. The goal of this set up was to explore how users would configure their workplace prior to interaction, without any influences from predefined token positions.

Measures: Measures were obtained from video recordings of participants interacting with the tabletop. The analysis was formalized and simple, objective data extracted. To gain insight into how effective were different token structures and the system repositories in supporting embodied thought, two traditional performance metrics were analysed:

- *Search time*: The mean duration between holding tokens. This is the time from when a participant places a token down until they pick one up again.

- *Hold time:* The mean duration a token is in the hand of a user, measured from pick up to drop down. This included time spent moving the token.

The reasoning behind such measures is that they enable the assessment of which offline epistemic strategies account for lower times to either: find the appropriate token from the system's repositories; or decide what is the appropriate course of action after picking up a token. These measures will also help determine if the strategies that most commonly and naturally occur are actually the strategies that allow users to perform quicker and more efficiently. To reduce variance, this data was filtered using a selection criterion: if users were judged to be distracted (e.g. starting a discussion with another user) or would start another action while holding or searching for a token, the measurement was discarded. Finally, after the completion of each

of the three tasks the experimenter would take note of the total number of tokens used in the final routine solution.

Qualitative data was also inferred from the videos, as to explore the strategies developed by each participant (or pair of participants) when organizing or moving tokens from the repository areas to the *sensable* surface areas of the application (see Figure 6.3). In the collaborative condition, verbal and physical communications between participants were also noted (e.g. pointing, touching, guiding or token passing) in order to establish whether participants were adopting distinctive roles (as discussed by Shaer et al. [SSV+11]). Finally, subjective data was also captured. All participants completed a short questionnaire that measured their engagement, affect and fulfilment while using the system [LGE08]. The questionnaire was 12 items long; four questions represented each of these three constructs. All questions were measured on seven-point Likert scales, ranging from strongly disagree to strongly agree. A brief semi-structured interview followed the study and asked participants to expand on their strategies for managing tokens outside of the system's sensing capabilities.

Results

In this section general data from both quantitative and qualitative measurements is presented. This is done to better understand what strategies users adopted when interacting with the system in term of: actions undertaken; number of tokens used (in both the repositories or in the routine area); bespoke repositories; and general engagement when using the application. Table 6.1 shows the basic experimental results.

The first finding is that, for single users, the duration of actions stays relatively steady or decreases as they become more familiar with the system

(search time: F (2) = 3.71, p = 0.05; hold time: F (2) = 2.98, p = 0.07), whereas the number of tokens used remained fairly constant (Task 1: 7.83/2.73; Task 2: 8.08/3.68; Task 3: 7.50/1.88). Despite previous work showing that experienced users are more likely to engage in epistemic actions [KM94], these results point out that it is not always the case, with factors such as task complexity or time pressure influencing in how much cognitive aid will users look for in tokens in the periphery of the interaction. Additionally, and while there was no significant differences between the duration of actions performed by single and paired users, the latter group used more tokens than the former (Task 1: $t(22) = -5.83$, p = 0.001; Task 3: $t(22) = -9.05$, p = 0.001). Lastly, single users were also faster than paired users in searching for tokens ($t(21) = -3.56$, p = 0.01) but also took longer to decide what to do with them (hold time).

Strategies: The strategies employed by users during task completion provide valuable insights into their cognitive activities. By managing tokens outside of the interaction space, users create a spatial structure where they know what to expect, where to look and what to look for [P95] – effectively lowering task complexity. Three distinct strategies of interaction emerged from the observation of the participants tackling each of the different tasks. The first strategy, *repository-focus* (**R**), was characterized by users who relied heavily on the repositories to manage tokens throughout the task. Such users would pick

	Task 1	Task 2	Task 3
Search time	Single: 2.87 (1.40) Pair: 4.49 (1.73)	Single: 2.37 (1.06) Pair: 4.57 (2.55)	Single: 3.06 (1.90) Pair: 5.25 (2.48)
Hold time	Single: 2.03 (0.73) Pair: 2.37 (1.65)	Single: 2.47 (1.15) Pair: 2.19 (1.26)	Single: 1.54 (0.44) Pair: 1.66 (0.47)
Number of tokens used	Single: 7.83 (2.74) Pair: 13.50 (1.98)	Single: 8.08 (3.68) Pair: 8.80 (1.55)	Single: 7.50 (1.88) Pair: 13.50 (1.31)

Table 6.1. Mean times for search, hold and number of tokens used in each of routines created (one per task). Standard deviation in brackets (single and paired users).

a token from the repositories, explore the available options (via the options zone of the interface, see Figure 6.3), and, if acceptable, lay the token in the appropriate location in the routine area of the interface; or alternatively, either reject the token (or cue them for later exploration) by returning it to a repository area. In the second strategy, *action-focus* (**A**), participants would move tokens directly from the repositories to the routine area, laying out routines from ideas established when the tokens were in the repositories. In the end, they would pick each token from the routine area and adjust the options in turn. The last strategy, *mixed-focus* (**M**), merged both R and A strategies. Additionally, participants in the distant condition adopted three unexpected repositories to manage their tokens (see Figure 6.4):

- *Hand* (**H**): participants would manage tokens in their hands and arms, sometimes creating a queue system to simplify token access and choice;

- *Interaction space* (**IS**): participants would manage their tokens directly in the tabletop's routine area;

- *Periphery* (**P**): participants would use either the edges of the tabletop or repository areas built into the interface to manage their tokens. In one particular case, a participant placed a nearby office chair by the system so that he could better manage his tokens there.

Figure 6.4. Three repositories that naturally emerged from participants in the distant condition (from left to right): hand (H), interactive space (IS) and periphery (P).

As participants in the distant condition become more proficient with the system they converged on strategy **A** and repository **H** (7 out of 12). The same was observed for participants in the local condition (who were primed by the setup to use the repository areas to hold their unused tokens), where 8 out 12 pairs used strategy **A** in Task 3. For these last participants, strategy **M** was the most popular for the first two tasks, with strategy **R** becoming less popular as users become more experienced with the system (it was not used in Task 3 at all). Additionally, participants who did not change their strategies were significantly more engaged than those who did (t (10) = -3.34, p = 0.01), and had significantly higher affect (t (10) = -3.42, p = 0.01). Lastly, participants in the distant condition opted for the hand repository in all of the tasks (Task 1: used by half of the participants; Task 2: used by 9 out of 12 participants; Task 3: used by 7 out of 12 participants).

Discussion

Klemmer states "we possess nearly unlimited modes of interaction with the physical world" [KHT06]. Through tangible interaction, users are now able to interact with digital resources through tools that are available to them in the real world. But human behaviour is complex and subject to a wide range of influencing factors [PW92]. As such, and as a result of a continuous interaction with the surrounding environment, users develop their own strategies to deal with the problems at hand. But due to the array of available options at any point during interaction, such strategies might not even be the most 'efficient' in objective terms [C01]. It is clear that proving users with such openness in terms of how they interpret their interaction with technology creates a new challenge for developers [FTJ08]. This disseration argues that this is one of the biggest challenges when designing tangible applications – how to accommodate not only task oriented actions, but also offline epistemic activities that lower the cognitive demands of the task at

hand and improve the general user experience. This challenge extends beyond design to include evaluation of such systems, as offline activities are both poorly defined and hard to predict – i.e., there is little literature suggesting what offline activities to look for [FTJ08]. This ultimately makes these activities hard to formally capture and represent using quantitative methods from HCI [PW92].

Addressing this lack, this chapter presents initial work exploring the design of interactive tabletop systems to support offline epistemic activities, especially those that aimed at offloading cognition into the environment to decrease task complexity. Although the study presented focuses on a particular task, the insights gathered on offline activities are intended to be broad enough to be useful and applicable to a wide range of future designs. As such, unlike traditional guidelines in HCI that focus on specific interaction mechanisms [e.g. MBS92, NL95], what is proposed are guidelines to support a broader set of activities that are relevant to the rapidly growing paradigm of tangible interaction. These guidelines are presented in the remainder of this section covering four different areas of interaction design: the form of repositories for managing tokens outside of a system's sensing capabilities; the size of these repositories; and their location in a tabletop system.

Repository form: The study manipulated the form of the available token repositories: a single, large, differentiated area versus four smaller, meaningfully categorized zones. Such distinctions appear in the literature (for example in the Marble Answering Machine's [A99] use of a single large cache for message-beads integrated into the machine itself in conjunction with individually labelled dishes on which to stash and aggregate specific groups of beads) and a key goal of the presented study was to explore how users take advantage of such spaces. The results indicate the zones were largely ignored. In 61% of the tasks in the distant condition, where participants were required

to bring relevant objects to the table, the arms and hands were used as the primary storage space for the tokens. Similarly, in the local condition where tokens were appropriately set up in the repositories prior to the study, participants rarely used the labelling on the zones to support their tasks. In contrast, they grouped or stacked the tokens according to their own schemes or mental structures [A12] (typically based on token similarity) during breaks between tasks.

There are a number of possible accounts for this behaviour. Explaining the tendency to hold tokens in their hands, several participants reported uncertainty about the function of the repositories – they formed a part of the tabletop (e.g. in that the graphical contents were digitally projected) and the behaviour of the system to tokens in the repository was perceived to be ambiguous. A clear recommendation is therefore that repositories need be clearly marked and possibly even visibly beyond the active sensing and display areas of a tangible system. When tokens were already present in the repositories users were much more likely to take advantage of these zones suggesting that such clear examples are very useful in communicating appropriate behaviours to novice users of tangible systems. The use of physical affordances (e.g. fixtures such as racks that are clearly designed to contain and store tokens [e.g. 2]) may be a useful mechanism by which to achieve this. However, despite using the zones, participants typically ignored the categorical structures in place in the four-zone condition in favour of their own schemes. This observation supports Dourish's notions that customization and appropriation are key to offline activity [D01] and suggests that the design of repositories might be best directed towards supporting these forms of personalization as opposed to pre-determined structures or categories.

Repository size: Users relying on the repository areas of the interface were observed de-stacking tokens whenever space was available, attempting to

ensure that the greatest number of tokens was always visible simultaneously. Furthermore, although participants in the local condition (where most used the hand repository) were faster than participants in the distant condition (where all used the repository areas) at locating desired tokens, the latter created more elaborate and complex solutions. This comes as no surprise, as people naturally seek to augment their cognitive abilities by, for example, reorganizing their environment [WVV11]. Spatial rearrangement aids problem solving by changing how users look at a problem, making objects within the task more perceptible [HB06]. This can be further explained through what Kirsh introduced as projection [K09b]. As mentioned in chapter three, projection is cognitive work that mixes perception (which relies on sensed information) with imagination (which relies on cognitive information) to yield mental augmentations that are anchored to physical artefacts. As such, knowing how much space should be reserved for offline activities is important, as it influences how many tokens are perceivable by the users which in turn impacts on how they think and plan for the task at hand.

However, as we have seen before in chapter two, a substantial challenge of tangible interaction is that physical tokens inevitably create clutter [SH10], so repository size affects how flexibly this can be contained. A consideration of the activities that repositories need support can therefore provide insights into the size they should adopt. The results for this study indicate that the space for offline activities needs to be big enough to support the highest number of visible tokens possible. Designers need consider both the number of tokens their system accommodates and the nature of the tasks users will be performing. For example, many board games (e.g. chess, checkers) may require the use and management of only one token at a time.

Repository location: Another important issue when considering repositories for unused tokens is in which area of the interactive tabletop display to place

them. One of the hallmarks of tangible interaction has been its capacity to better support collaboration, e.g. as tokens provide users with multiple access points that lower the threshold for participation [HSC+09] and thus are understood as resources for shared activity [FTJ08].

In terms of collaboration on a single task, paired users were observed taking two distinct roles (sometimes interchangeable during a task): one would act on the tabletop, laying tokens in their final positions and deciding on which options to set-up; while the other would take a supporting role, by choosing the next token to be used, providing verbal or visual tips (e.g. pointing) or discussing several possibilities and scenarios. This an interaction profile previously identified by Shaer et al. [SSV+11] as the driver-navigator (see Figure 6.5). And while it is the physicality of the interaction that allows for a clear understanding between users [KHT06] and a distinct separation of roles, the current idea of multiple users interacting simultaneously and synergistically on a tabletop – one of the hallmarks of tangible interaction – is dislocated from a real-world setting. As such, the search for an ideal location for areas to support offline activities does not need to be made in ways that make them perfectly accessible to every user in the system. Moreover, as it was shown that both experienced single and paired users tend to converge to a common strategy (**A**), conducting preliminary studies over prototype systems seems like an ideal way to reveal possible candidate locations for areas that support offline activities.

Figure 6.5. Two participants engaging with the system in what Shaer et al. [SSV+11] describes as the driver-navigator collaboration profile.

In conclusion, this chapter has demonstrated that a small change in an interactive tabletop's size or shape can have a dramatic impact on its use

[TT06]. While not comprehensively answering why this is the case, this work contributes to the field of tangible interaction by suggesting that changes in support for offline epistemic activities may be partly responsible for this. Indeed, such activities are a unique aspect of tangible systems and only by examining such distinctive properties will we be able to elucidate the real benefits and value of the tangible paradigm.

EPISTEMIC ACTIONS AS RESEARCH TOOL

7 This dissertation has observed and argued so far that while users and designers continue to be drawn by the allure of physically handling digital data, it remains challenging to understand and quantify the genuine benefits of tangible interaction [ZAM+12]. Indeed, this dissertation argues there is currently no systemic account of the underlying properties or qualities of tangible systems that can explain or justify their enduring appeal. There is no comprehensive way to answer questions regarding the true value provided by tangible systems to its users. However, steps are being taken to develop such answers and explanations. As mentioned in chapter three, one fertile source is the cognitive science literature that focuses on embodied (or situated) cognition – especially **epistemic actions**. In the field of HCI, these kinds of activities are comprehensively documented in, for example, work analysis of air-traffic control – an activity in which operators rely heavily on physical paper strips to mediate their complex and safety-critical work tasks [M99]. Other authors have depicted how users leverage external (non-mental) structures as tools to simplify cognitive work [K97]. For example, experienced jigsaw puzzle-solvers often cluster physical pieces together (e.g. by colour) to simplify subsequent processes of visual search and recognition [AW13]. These kinds of account are important as they provide a basis for explaining the appeal and value of tangible systems – they cast light on the ways these really

provide benefits to their users, potentially steering future design and development efforts.

#	**Epistemic actions**
A1	**Manipulation of an artifact**
A2	Spatial arrangement of artifacts in relation to one another, the task environment, or the users
2.1	Cluster or group artifacts together
2.2	Divide workspace into several stations in which only a subset of actions are afforded
2.3	Place an artifact in a contrasting environment
2.4	Rearrange a representation
2.5	Clear and clean clutter
A3	Parallel use of two artifacts, two representations, or an artifact and a representation
A4	Artifact trial-and-error positioning
A5	Shuffle artifacts
A6	Compare an artifact with a possible destination or other artifacts
A7	Mark an artifact
A8	Test the state or response of a system, model or other user
	Manipulation of a tool
T9	Tag or annotate an artifact
T10	General notes and annotations
T11	Use of a tool to physically constraint the user or the use of other artifacts and tools
T12	Build a model or external representation
	Bodily action
B13	Use the body to externalize an internal process
B14	Talk or gesture to guide and direct attention
B15	Move the body, problem space, or representation

Table 7.1. A list of all the 20 types of epistemic actions present in the ATB video-coding framework, which groups actions by those performed with an artifact (an epistemic action in itself), a tool, or the user's body. Five types of epistemic action are grouped under A2, a broad type of action.

However, applying principles from cognitive science to the design of tangible systems remains a challenging and intricate task. Although work on this topic remains embryonic, several distinct approaches exist. One key thread, instantiated as illustrative design frameworks [e.g. HB06, JGH+08, KHT06], aims to provide high level guidance and recommendations for how tangible systems can best be created. A second strand, more directly related to the literature on embodied cognition, seeks to expand our repertoire of metrics for understanding and assessing performance with tangible systems. Basically, it argues that theoretically grounded techniques that enable us to rigorously and empirically examine the mechanisms by which users rely on physical objects to aid their cognitive work will help us assess (and ultimately learn how best to design) such systems. This latter approach has been explicitly explored in the context of Kirsh et al.'s [KM94] categorization of behaviours into either pragmatic or epistemic actions. While Kirsh originally explored this idea in the context of a purely virtual computer game (Tetris), Antle et al. [AW13] were arguably the first to apply it to a tangible system. Their work compares tangible and touch interfaces to a classic problem-solving task – a jigsaw puzzle – and they use video-coding analysis to classify user behaviours as either pragmatic or epistemic. Their results suggest that the tangible representation afforded more natural and efficient epistemic strategies such as clustering pieces to improve subsequent visual search, or relying on the elevated edges of the table to help structure the puzzle.

While this work provides evidence for the conjecture that tangible interaction aids performance of epistemic actions, one major weakness is in terms of the granularity with which the activities are recorded (as illustrated in chapter five). Basically, work that defines and discusses epistemic activity is typically highly specific and contextual – particular epistemic behaviors are described as being used to achieve particular ends in particular situations. However, the evaluation frameworks in HCI are broad and general, often reducing the

diversity of epistemic activity to a single categorical label. This chapter aims to address this issue by expanding an existing video-coding framework [A12, AW13] that categorizes hand actions to include a detailed classification scheme for epistemic activity – the ATB (Artifact, Tool and Body) framework. This chapter argues that this framework will contribute to our understanding of how epistemic actions are used in relevant tasks, providing researchers with a tool to more systematically assess this complex type of behavior in tangible interaction.

In terms of HCI, this tool has two objectives. Firstly, it is intended as a mechanism to evaluate tangible systems in terms of the type, diversity and appropriateness of the epistemic actions they support, and in terms of the impact these actions can have on more traditional metrics such as performance time or errors. Secondly, in the long term, this chapter argues that a series of such evaluations will result in a corpus of knowledge describing the use of epistemic actions in real tasks. This data can be used as the basis for grounded, practical design knowledge on how to create novel systems that truly support epistemic actions, and thus, improve our ability to design tangible interaction that is natural and meets the real needs of the user. As such, this chapter makes two contributions. Firstly, it describes a detailed framework of epistemic activity based on a systematic literature review of 78 papers. Although loosely related prior classifications exist [e.g. K95b, KM94], the framework presented in this chapter is the first to be based on a systematic review, the first to aim for a focused, fine-grained description of epistemic behaviour and the first to be specifically directed towards the development of an actionable empirical tool for capturing and expressing observed epistemic actions. Secondly, this chapter presents an initial experiment to explore the framework's reliability, validity and predictive power. This substantial lab study involved 60 participants across three countries completing a physical task – a jigsaw puzzle. Three raters analysed

the data to support commentary on reliability. Validity is explored by contrasting the video-coding results among our purposely diverse participant group with other measures such as spatial ability tests and task completion rates and times. The outcomes of this study provide insights into epistemic activity and demonstrate the usefulness of the framework as an analytic tool that other researchers can apply in their own design and evaluation activities in the field of tangible interaction.

The ATB Video-Coding Framework

The work in this chapter builds on the action classification framework presented by Antle et al. [A12, AW13]. In their work, an action can be classified as either a *direct placement* (**DP**), an *indirect placement* (**IP**) or as *exploratory* (**EXP**). In the puzzle task they studied, a DP action corresponded to those situations where users already know where to place a piece before picking it up, leading to a fast and direct transition between acquiring a piece, moving to the final destination and correctly placing it. IP represented similar outcomes but described situations in which users are not initially certain of where to position the pieces they pick up. As such, they translated or rotated the piece while searching for its correct destination. Finally, EXP represented those actions where pieces do not end in their final and correct position. As with IP, if these intermediary actions make the task easier for the user they are considered epistemic (e.g., if a user organizes pieces into different piles for subsequent identification and retrieval).

Antle's framework [A12, AW13] provides basic features that enable the study of epistemic actions. It allows researchers to measure epistemic activity levels within a task, reporting on the frequency, duration and moments at which epistemic actions occur. While valuable, this chapter argues that this classification is too broad to fully articulate and explain the role of epistemic

action in problem solving. For example, different types of epistemic action may be used in different tasks (as demonstrated in the previous chapter), and recording such variations in detail will better characterize the role and importance of epistemic activity. This chapter argues that only by considering epistemic actions at a fine-grained level of detail will we be able to understand not just how many epistemic actions are performed during a task, but also which epistemic actions are chosen and to which purpose. In terms of tangible interfaces, a detailed classification scheme will support investigations of what interface elements facilitate what epistemic actions, quantifying the differences between novel systems, and allowing designers to tailor interaction that better supports users' natural, epistemic behaviours.

Framework Development and Use: To develop the ATB framework an extensive literature review was conducted with the goal of capturing a wide range of epistemic activity descriptions. A set of keywords was used to conduct a literature search on both Google Scholar and Science Direct. The search terms were 'epistemic action(s)', 'complementary action(s)' and 'complementary strategies'. The first 60 results from each of these searches were kept for further inspection. Additionally, papers referencing seminal work in the area (specifically [K95] and [KM94]) and including the keywords defined above were also retained. Ultimately, 78 papers were obtained through this process – a typical number for meta-analysis papers in the area of HCI [e.g. HL07]. Each paper was then inspected for any mention of actions that could be interpreted as epistemic, or were directly treated as epistemic, and quotes such as: "(...) *preparing the workplace, for example, by partially sorting nuts and bolts before beginning an assembly task in order to reduce later search* (...)" [T85, p. 515] were extracted. These represented concrete examples of epistemic actions from research literature in a range of fields (such as mathematics, cognitive science, HCI and design) from the last three decades. A complete list of references for all 78 publications can be viewed in Annex 2.

In total, 335 quotes were compiled. With these, a pair of researchers highly familiar with the literature on embodied cognition worked collaboratively to create an affinity diagram that identified different clusters of epistemic actions. Quotes judged to depict actions with unclear epistemic value were discarded. This process led to the identification of 20 types of epistemic action based on a subset of 225 of the original quotes. These were then grouped by actions performed with 1) task artifacts (e.g. objects marked with fiducials), 2) tools (e.g. a pencil that can be used for annotations) or 3) the users own bodies (summarized in Table 7.1). Annex 3 contains a full scheme describing these 20 types of epistemic action.

These categories are then used as the basis for classifying behaviours through video-analysis, according to the following procedure. Firstly, raters should categorize actions as being either DP, IP or EXP, as in Antle et al.'s framework [A12, AW13]. After this process is completed, raters should review each action that can contain epistemic activity (i.e. those coded as either IP or EXP) and match these to one (or more) of the 20 types of epistemic actions identified. For a graphical workflow of how to video-code using the ATB framework, please consult Annex 4. A coding scheme file was created to facilitate the process of video coding with Anvil, a free and popular video-coding tool. This file can be accessed in Annex 5.

Applying the ATB Framework: An Initial Study

To determine the usefulness of the ATB framework in capturing and distinguishing among different epistemic actions, and the fundamental value and worth of this kind of information, this chapter presents an observational study of users performing a classical problem-solving task – a jigsaw puzzle. This task was selected as there is a large body of work on epistemic actions

using puzzles in both HCI [e.g., A12, AW13], and cognitive science [e.g. 15, 16, 17] [e.g., K95, K96, K99]. Furthermore, puzzle metaphors are commonplace in the design of tangible systems [e.g., HJ07]. The goal of this initial study was to explore the reliability and sensitivity of the framework, and to assess its internal and external validity, and predictive power. To meet these classic methodological objectives, a diverse participant pool was recruited (from Korea, Canada and the Netherlands) and a range of spatial and subjective workload tests were performed to establish the main results in a theoretical context.

Experimental Design and Participants: All participants in this study completed a single condition. In total, there were 60 participants, 20 of whom were Korean, 20 Dutch and 20 Canadian – 10 male and 10 female participants from each nationality. The study also took place at three sites, one in each of these countries, with all participants residing in their respective countries of origin. Participant ages ranged from 20 to 76 (M = 27.33, SD = 10.72) and occupied a wide range of professions, from undergraduate to postgraduate students, to sailors, game designers, artists, drivers and writers. Before the study, all participants filled in a brief online questionnaire to exclude puzzle hobbyists. To motivate participants to perform to the best of their capabilities, a prize of $25 (or equivalent) was awarded to the fastest participant to solve the puzzle in each of the three countries. In addition to this prize, Korean and Canadian participants received a $10 compensation for participating. Dutch participation was not compensated due to different funding policies in the three research groups.

Procedure: Each session involved a single participant solving two puzzles, and performing an additional test at the beginning and end of the study. All tasks were performed in small and otherwise empty offices, and all sessions followed the same structure. Sessions commenced with a brief introduction to

the first task, a paper folding spatial ability test [EFH+76]. Participants were then introduced to the first jigsaw puzzle, an unmeasured practice task, which they were asked to solve in a maximum time of 10 minutes. This was followed by the main task of the study, a second jigsaw puzzle which participants were asked to solve in a maximum time of 15 minutes. Though the two puzzles presented different images, each consisted of 70 pieces and was 38x26cm in size. The order in which the puzzles were presented was the same for each participant. Before starting the main puzzle participants were reminded of the monetary prize. During both puzzle tasks participants were left alone to ensure that their epistemic actions were unmediated and private. Both tasks were recorded on video for later analysis. At the end of the main puzzle, participants completed a subjective test to measure the perceived workload of the main task.

Measures: In addition to recording the time that it took participants to finish the puzzles, the following metrics were used:

- *Spatial ability (paper folding test)* – Upon starting the study, participants were required to solve two sets of spatial tests [EFH+76], each in under three minutes.

- *Video-coding framework* – Several metrics were derived from the data obtained through the video-coding framework being examined. These include the mean number of pragmatic (coded as **DP**) and epistemic actions (coded as either **IP** or **EXP**) performed; and, the individual mean frequency of each of the 20 types of epistemic action in the framework. These frequency metrics include both aggregate and running means, on a minute-by-minute basis.

- *Subjective Workload* – Each participant completed the NASA TLX [HS88] after the main puzzle task.

Results and Discussion

The study presented in this chapter has two goals. Firstly, to assess the usefulness and correctness of the ATB framework. Secondly, to introduce epistemic actions as a comprehensive new performance metric for systems incorporating tangible interaction. To facilitate this discussion, the results of the study are divided into four sections: *spatial ability*; *framework reliability*; *framework validity*; and *framework predictive power*. Workload test results are described when relevant throughout this section.

Spatial Ability: The paper folding test [EFH+76] was performed to provide meaningful data on participant's inherent spatial ability, as it could have a significant impact in how they perform during a jigsaw puzzle task (and how they might ultimately rely on epistemic actions). As expected, participants that were able to finish the main task in less than 15 minutes obtained higher scores than those who did not (see Table 7.2). An independent-samples t-test

Participant group		Spatial ability	Overall Workload	Overall time to finish
Puzzle completion	**Finished** **Incomplete**	15.11 (4.14) 11.67 (6.27)	8.74 (2.31) 11.36 (2.75)	9:38.10 -
Gender	**Male** **Female**	15.61 (3.59) 13.13 (5.62)	9.36 (2.69) 9.19 (2.48)	10:49.22 10:46.26
Cultural background	**Korean** **Dutch** **Canadian**	17.48 (2.76) 15.06 (3.50) 10.56 (5.19)	8.79 (2.47) 10.29 (1.57) 9.00 (3.24)	12:44.00 08:58.54 10:40.48
Rater	**Rater 1** **Rater 2** **Rater 3**	15.15 (4.70) 13.94 (5.07) 14.01 (4.89)	9.18 (2.56) 9.96 (2.27) 8.67 (2.76)	10:52.03 11:34.39 09:57.00

Table 7.2. Mean scores for the spatial ability paper folding test (higher is better), the NASA TLX (lower is better), and the overall time to finish the task (mm:ss.ms). Std. dev. in brackets.

revealed that this difference was not statistically significant ($p = 0.062$), which may be attributable to the disparity between the two group sizes – of 60 participants, only 13 failed to complete the task. Additionally, each group of twenty participants coded by the three individual raters exhibited statistically similar spatial scores (one-way ANOVA, $p = 0.744$), demonstrating the equivalence of these groupings. These results will be further discussed in the following sections.

Framework Reliability: The first step when introducing a new measurement instrument, such as a video-coding framework, is to establish its reliability. To do this, three independent raters applied the video-coding framework proposed in this paper to the videos of the 60 participants performing the main puzzle task. Each rater coded 20 videos: 10 male and 10 female participants, with either six or seven of these from each of the three cultural backgrounds sampled. Events were classified with a timestamp, Antle's classification (DP, IP or EXP), and one or more types of epistemic action (if epistemic activity was observed). Additionally, two of these raters acted as second-coders for eight of the 60 videos, providing a selection of double-coded content. These two raters obtained a substantial agreement in terms of Antle et al.'s three categories of action (85.6%), and in the main three categories of the ATB framework – artifact, tool and body (75.7%). These high level results illustrate how the ATB framework can inform how future tangible systems are designed and implemented. Basically, these results demonstrate that the framework enables researchers to reliably identify the most common source of epistemic activity during interaction with a system, be it through artifact, tool or body. Understanding where the focus of epistemic activity lies will help direct research and design efforts in the most appropriate directions. For instance, if users depend on task artifacts (as we observed in our task, see e.g., Figure 7.2), attention can be directed to improving these. Alternatively, if users look at tools to alleviate cognitive

burdens, these should serve as inspiration for new artifacts or systems that effortlessly accommodate tool use. Finally, if a task is most suited for bodily actions, designers can focus on how best to sense and enable these.

Beyond this point, the strength of the ATB framework is in the granularity of the epistemic actions recorded. In regards to its 20 different types of action, coders reached a moderate level of agreement of 60.9%. A confusion matrix revealed that 87.86% of misclassified actions belonged to two groups: actions in which one rater observes some epistemic activity, but the other does not (61.15% of all misclassifications) and 'vertical' misclassifications, where both raters use nested epistemic actions to classify a particular event (e.g. between **IP**: **A2** and **IP**: **A2.2**, 26.7% of all misclassifications). The most common of these misclassifications occurred with action **A1**, under which all epistemic actions performed with an artifact fit.

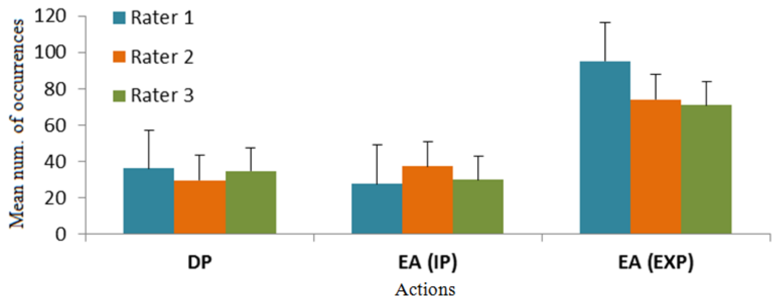

Figure 7.1. Mean number of pragmatic (DP) and epistemic actions (EA, grouped by actions classified as IP or EXP). Overall data grouped by rater, standard deviation in bars.

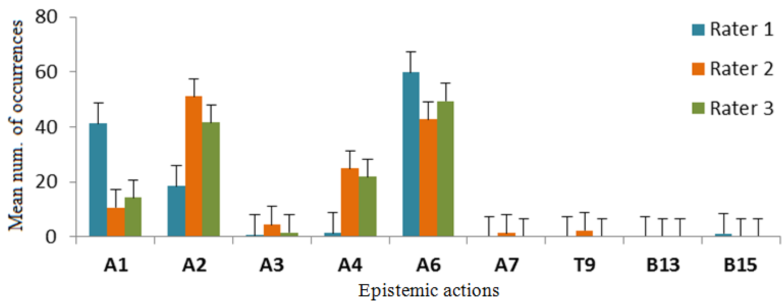

Figure 7.2. Mean number of individual types of epistemic action, with A2 grouping actions classified from 2.1 to 2.5. Data clustered by rater, standard deviation in bars.

Given the limited number of videos double-coded by two raters, additional statistical tests were performed on data from all the 60 sessions to further explore and qualify agreement levels between the three individual raters (subsequently termed *R1*, *R2*, and *R3*) in terms of mean categorical response rates across all rated data. This data is summarized in Figures 7.1 and 7.2. A one-way ANOVA was used to study the differences between these datasets – Welsch's F and Games-Howell post-hoc tests were used when the assumption of homogeneity was violated. All three raters reported similar mean numbers of epistemic actions coded as **IP** ($p = 0.090$), and R2 and R3 reported similar mean numbers of actions coded as **EXP** ($p = 0.952$), **A1** ($p = 0.343$), **A2** ($p = 0.573$), **A4** ($p = 0.699$) and **A6** ($p = 0.448$) – the four most common epistemic actions coded by these raters. While these four epistemic actions were also the most commonly reported by R1, most of their occurrence rates significantly varied from the other two raters ($p = 0.015$). Close examination of this pattern suggests it can be explained by the confusion matrix described earlier – R1 tended to classify events with the broad **A1** action (Figure 7.2).

Taken together, these are promising results that vouch for the reliability of the ATB framework as an instrument to record epistemic work. Substantial to moderate agreement levels were attained on different levels of the framework and examination of the raw data shows clear parallels between rater performance. Observed misclassifications fall in a limited number of acceptable types. This chapter argues this data effectively illustrates the reliability of the framework.

Framework Validity: In this section the *external validity* of the ATB framework is assessed – the generalizability of the framework to a broad participant group. Two different methods are used to achieve this: (1) contrasting the obtained results with current theory on gender differences in spatial ability;

and (2), comparing the results of participants from different cultural backgrounds.

There is a long tradition of studying of cognitive differences between the genders, with predictably conflicting and controversial results. While men are often regarded as having higher spatial ability, some studies suggest gender differences are small [e.g., H81] and, indeed, diminishing [VVB95]. Our own results show that male participants obtained higher spatial scores than female participants (p = 0.029), but both reported a similar perceived workload when performing the task (p = 0.797) and finished with statistically similar mean times (p = 0.893) – see Table 7.2. This chapter argues that this can be explained by examining the results obtained with the ATB framework. These show no statistically significant differences between the mean number of actions performed between the genders (see Figure 7.3 and 7.4): **DP** (p = 0.908), **IP** (p = 0.728), and **EXP** (p = 0.840); **A1** (p = 0.801), **A2** (p = 0.244), **A4** (p = 0.985), and **A6** (p = 0.842). This chapter suggests that by performing the same

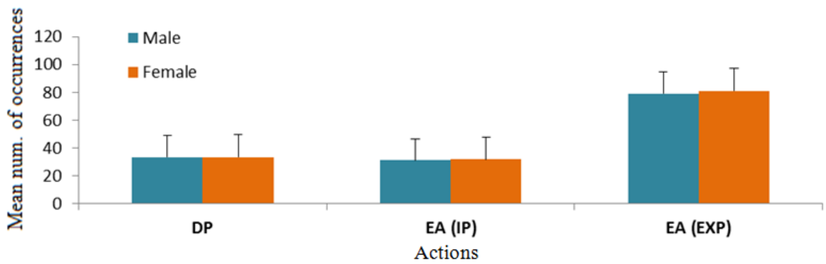

Figure 7.3. Mean number of pragmatic (DP) and epistemic actions (EA, grouped by actions classified as IP or EXP). Overall data grouped by gender, standard deviation in bars.

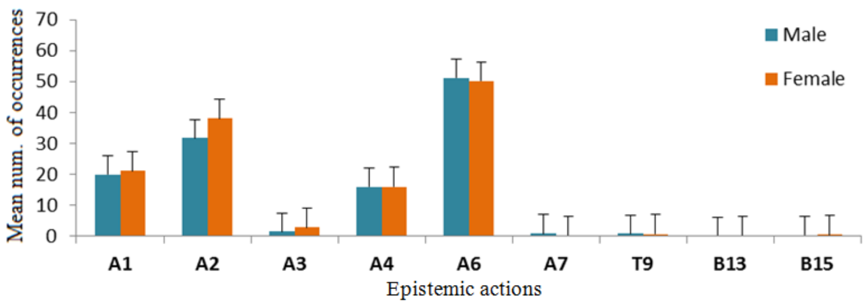

Figure 7.4. Mean number of individual types of epistemic action, with A2 grouping actions classified from 2.1 to 2.5. Data clustered by gender, standard deviation in bars.

number of epistemic actions as their male counterparts, female participants were able to make up for any differences in spatial ability, and indeed that adopting appropriate epistemic behaviours may be more important in this task than high spatial ability.

Furthermore, the experimental data also supports the framework's ability to generate coherent results across broad and varied participant groups. Specifically, the spatial tests recorded a lower spatial score from Canadian participants when compared to both Korean and Dutch participants (one-way ANOVA, Games-Howell post-hoc tests: $p = 0.001$ and $p = 0.003$, respectively). As with the gender groups, however, Canadian participants reported similar workload levels as participants from the other two cultural backgrounds (one-way ANOVA, $p = 0.055$) – see Table 7.2. This chapter argues this is attributable to the similar mean number of epistemic actions performed by participants from each cultural group: **A1** ($p = 0.145$), **A2** ($p = 0.064$), **A4** (p =

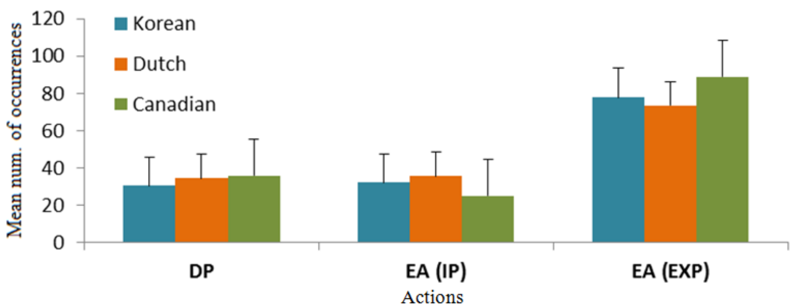

Figure 7.5. Mean number of pragmatic (DP) and epistemic actions (EA, grouped by actions classified as IP or EXP). Overall data grouped by country, standard deviation in bars.

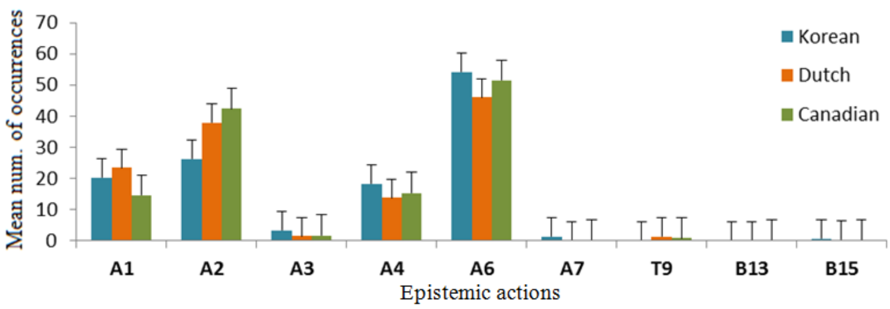

Figure 7.6. Mean number of individual types of epistemic action (A2 groups actions from 2.1 to 2.5). Data clustered by country, standard deviation in bars.

0.561) and **A6** (p = 0.329). More so, Canadian participants completed the main task in an average of 10 minutes and 40 seconds, two minutes and four seconds quicker than Korean participants (independent samples t-test, p = 0.037). The reason for this result may be in the first minute of the task [AW13, K97], where Canadian participants performed almost twice the number of epistemic actions than Korean participants. This idea will be explored in more detail in the next section.

Framework Predictive Power: This section explores whether the ATB framework records data that meaningfully relates to other performance metrics (extending the discussion of validity), and whether or not it offers novel explanatory insights into participants' epistemic work. It does this by contrasting the results of participants who successfully finished the main

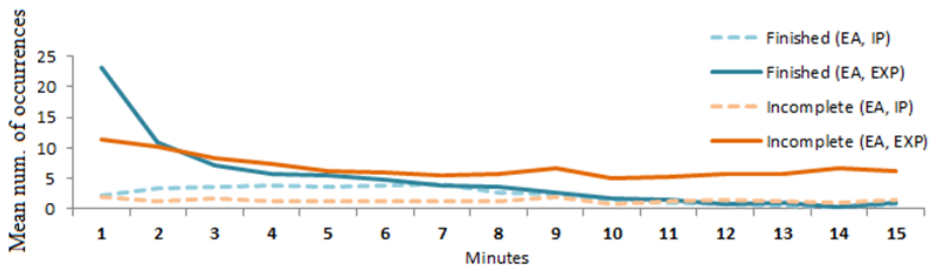

Figure 7.7. Running mean number of epistemic actions (EA, grouped by actions classified as IP or EXP). Overall data grouped by puzzle completion.

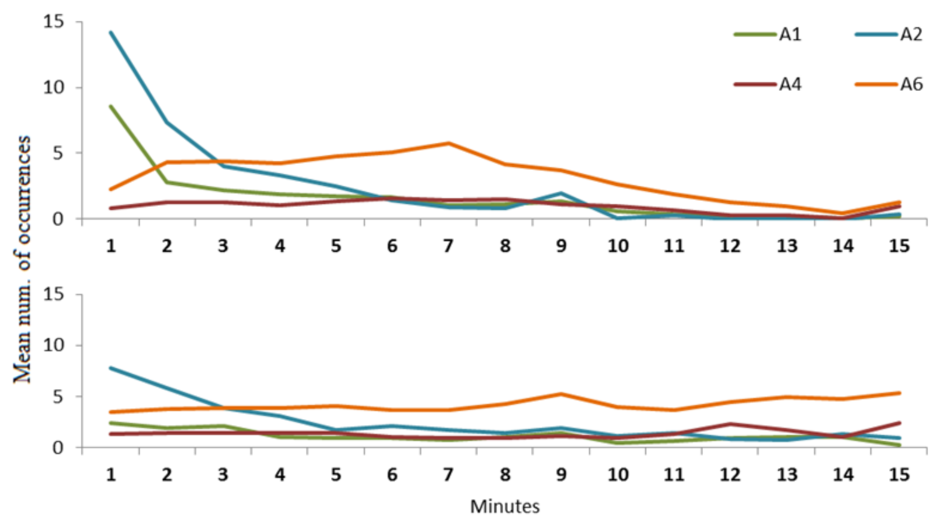

Figure 7.8. Running mean number of four types of epistemic action (A2 groups actions from 2.1 to 2.5). Data clustered by puzzle completion: finished (top) and incomplete (bottom).

puzzle task with those who did not. Specifically, informed by prior suggestions that early performance of epistemic actions is important in successful performance of spatial problems such as puzzle tasks [e.g., AW13, K97], this chapter firstly examined the impact of the rate of epistemic action performance in the first moments of interaction (see Figures 7.7 and 7.8). A linear regression showed that the aggregate number of epistemic actions performed in the first minute of the task were statistically significant in predicting participant's required time to finish the task (adjusted $R^2 = 0.177$, $F(1, 58) = 12.501$, $p = 0.001$). These results reinforce current suggestions [e.g., AW13, K97] that successful use of epistemic actions relates not simply to frequency of activity, but in knowing when it is worth performing them.

Following is examination of this data at the level of individual epistemic action as performed over the entire experimental task. Figure 7.8 shows this data for the most common epistemic actions performed (Table 7.1: A1, A2, A4, and A6, accounting for 96.48% of all epistemic actions). A multiple regression revealed that the ocurrence of these actions in the first minute of the task was a significant predictor of the time it took participants to complete the puzzle (adjusted $R^2 = 0.232$, $F(4, 55) = 4.158$, $p = 0.005$). Interestingly, not all of these variables added statistically significantly to the prediction (see Table 7.3). As such, this chapter suggests the ATB framework was successful in identifying which epistemic actions are more relevant and helpful for the user in the context of the early stages of a puzzle task (**A1** and **A2**), which were not

Variable	*B*	*p*	*Beta*
Intercept	12.663	-	-
A1	-0.077	0.008	-0.378
A2	-0.104	0.001	-0.470
A4	0.139	0.645	0.058
A6	-0.101	0.393	-0.110

Table 7.3. Regression coefficients and *p* values of a multiple regression that successfully predicts the time it takes participants to finish the main puzzle by measuring how many EA (A1, A2, A4, and A6) occur in the first minute of the task.

particularly helpful (**A6**), and which had (non-significantly) detrimental effects on user performance (**A4**). These findings provide evidence that ATB framework is a useful tool that can highlight what epistemic work is suitable for what kinds of problem and, this chapter argues, this kind of knowledge is valuable for the both the design and assessment of tangible systems.

Future Work and Conclusions

This chapter provides the basis for studies capturing granular data about epistemic activity. It does so by presenting and validating a novel framework that enables the measurement and detailed categorization of epistemic actions. Not only did the presented study test methodological aspects of the framework, it also took concrete steps towards developing evaluation metrics specifically targeted towards user experience in tangible systems. Limitations of this work include that, while it presents a study that looks at a substantial and diverse group of participants, it focuses on a single spatial task. While the framework was informed by epistemic actions collected from various fields, and thus should be applicable to most problems with physical properties, the most pressing future work lies on applying the framework to additional tasks. This work will allow us to reinforce, extend and generalize the findings currently presented. These include understanding the importance of actions **A1** (*Manipulation of an artifact*) and **A2** (*Spatial arrangement of artifacts in relation to one another, the task environment, or the users*) in other spatial problems and developing knowledge about which epistemic actions are relevant for non-spatial tasks. In combination with the work presented in chapters five and six, these understandings will ultimately lead to design guidelines specific to tangible interaction, allowing for the development of new interactive systems that support not only goal-directed, pragmatic actions, but epistemic strategies that enable users to apply natural, real-world knowledge to interaction with digital information.

CONCLUSION

8 The work presented in this dissertation focused on a relevant issue for modern HCI: the lack of paradigm-specific evaluation tools that can inform the design of future computer systems. Current techniques from the HCI toolbox were tailored for GUIs and thus, focus on evaluating user performance by looking at task-centric actions. As such, their use in the evaluation of new interaction paradigms such as physical or ubiquitous computing is challenging, as these expand computation to the users' real-world surrounding and thus support much broader, open-ended actions and activities. This current issue in HCI is ultimately illustrated by a lack of quantitative validation of a wide range of benefits that were initially attributed to these novel interaction paradigms.

To address this limitation, this dissertation focused on a novel and popular interaction paradigm in HCI, tangible interaction. It then drew from the literature on embodied cognition, a novel perspective in cognitive science, to explore different ways in which tangible systems could be evaluated. To frame this exploration, this dissertation focused on a particular characteristic of tangible, physical interfaces: *tangible thinking*, or how users effortlessly rely on the physical environment to lower the cognitive challenges of the task at hand. This framing lead to the identification and categorization of a rich set of

actionable metrics for the evaluation of such physical systems: epistemic actions. The potential of such metrics for HCI is explored throughout the three main chapters of this dissertation. These directly addressed the four research questions highlighted in the introductory chapter of this dissertation:

1. Can Kirsh's findings on Tetris players be replicated in other digital systems? Do epistemic actions increase user performance in those situations?

Kirsh et al. observed that the key for the success of experienced Tetris players was not in performing actions more quickly than inexperienced players, but in performing more epistemic actions. Similar results were observed with participants playing Four-in-a-row (chapter five) and solving jigsaw puzzles (chapter seven). Chapter five described the impact of epistemic actions in user performance by demonstrating that when all three competing participants performed a high number of epistemic actions the game was likely to end in a draw (a balanced game), and that when one participant performed more epistemic actions than his opponents he was most likely to win the game. Chapter seven explored this phenomenon in more detail, demonstrating the importance of performing specific epistemic actions at specific times during interaction – participants who finished the assigned jigsaw puzzle in less than 15 minutes where observed performing almost twice as much epistemic actions in the first minute of the task, compared to participants who were not able to complete the puzzle in the same amount of time.

2. Do different input modalities (e.g., physical tokens, touch, mouse) to a common computer system evoke different epistemic actions?

While the idea of *tangible thinking* and epistemic activity is normally associated with tangible interaction – and normally credited as a natural benefit of tangible systems – it was through a purely digital game that Kirsh

et al. introduced the concept of epistemic actions. The broad nature of such actions was illustrated in chapter five, as different input modalities to Four-in-a-row naturally supported different types of epistemic actions: participants on the mouse input condition favored the hovering feedback epistemic action; participants on the touch input condition favored pointing without a disk in hand; while participants in the tangible input condition performed epistemic actions quicker than in the other two conditions. These results highlighted not only how resourceful humans can be, even in the presence of purely digital tools, but also how challenging it is to measure, quantify and understand the benefits of tangible systems.

3. How do different design choices constraint or enable different epistemic behaviour?

Inspired by the intricate results obtained in chapter five, chapter six explored how different design choices impacted the epistemic strategies employed by users during task completion – especially actions performed outside of the sensing capabilities of the tangible system developed. Several observations were discussed, including how participants were likely to appropriate vacant or less important areas of the interface to perform specific epistemic actions, and how such participants were capable of producing more complex solutions than participants that were not so creative in the use of the space available. This chapter concluded with different design insights regarding the form, size and location for such areas where users are free to explore suitable epistemic strategies.

4. Can epistemic actions serve as metrics in novel evaluation tools for modern interaction paradigms in HCI?

Much like the first research question, the answer to this question is addressed in both chapters five and seven. Chapter five illustrated how the use of epistemic actions as metric for user performance led to a deeper understanding of the differences between graphical and tangible interaction paradigms. Chapter seven materialized this approach into a full-fledged video-coding framework that enables researchers to categorize epistemic actions into 20 different types of action. Chapter seven also described the results of a study validating this framework and demonstrating its ability to identify which epistemic actions positively affect user performance, and more importantly, when should these be performed to maximize their benefits.

In sum, the work presented expands our understanding of epistemic actions and suggests different ways of using these as metrics for novel computer systems. As such, it fulfils the three main goals of this dissertation: it provides new empirical knowledge attesting to the benefits of using epistemic actions as metric in tangible interaction; it explores different design solutions for the support of such actions; and, through the ATB framework, provides a novel and validated evaluation tool of epistemic actions for tangible systems. In addition, while the focus of this dissertation was on tangible interaction, the results presented should be useful to any researcher or developer working on modern interfaces that are instantiated in the real, physical world – e.g., ubiquitous computing, the internet of things, wearable technology. This idea is reinforced in chapter seven, where a physical but technology-neutral task was used to validate the framework presented. The ultimate goal of this framework is to enable the HCI community to generate a broad and novel body of knowledge of the different characteristics and benefits of future physical, computer-aided systems.

REFERENCES

[A79] Aish, R. (1979). 3D input for CAAD systems. *Computer-Aided Design*, *11* (2), 66-70.

[A99] Abrams, R. (1999). Adventures in tangible computing: The work of interaction designer 'Durrell Bishop'in context. *magistrska naloga, Royal College of Art, London*.

[A03] Anderson, M. L. (2003). Embodied cognition: A field guide. *Artificial intelligence*, 149 (1), 91-130.

[A07] Antle, A. N. (2007, February). The CTI framework: informing the design of tangible systems for children. In *Proceedings of the 1st international conference on Tangible and embedded interaction* (pp. 195-202). ACM.

[A12] Antle, A. N. (2013). Exploring how children use their hands to think: An embodied interactional analysis. *Behaviour & Information Technology*, *32* (9), 938-954.

[ACD09] Antle, A. N., Corness, G., & Droumeva, M. (2009). What the body knows: Exploring the benefits of embodied metaphors in hybrid physical digital environments. *Interacting with Computers*, *21* (1), 66-75.

[AKY00] Alibali, M. W., Kita, S., & Young, A. J. (2000). Gesture and the process of speech production: We think, therefore we gesture. *Language and cognitive processes*, *15* (6), 593-613.

[AN84] Aish, R., & Noakes, P. (1984). Architecture without numbers—CAAD based on a 3D modelling system. *Computer-Aided Design*, *16* (6), 321-328.

[AM00] Abowd, G. D., & Mynatt, E. D. (2000). Charting past, present, and future research in ubiquitous computing. *ACM Transactions on Computer-Human Interaction (TOCHI)*, *7*(1), 29-58.

[AW09] Aarts, E., & Wichert, R. (2009). *Ambient intelligence* (pp. 244-249). Springer Berlin Heidelberg.

[AW13] Antle, A. N., & Wang, S. (2013, February). Comparing motor-cognitive strategies for spatial problem solving with tangible and multi-touch interfaces. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction* (pp. 65-72). ACM.

[B96] Brooks, R. A. (1986). A robust layered control system for a mobile robot.*Robotics and Automation, IEEE Journal of*, *2* (1), 14-23.

[B00] Beaudouin-Lafon, M. (2000, April). Instrumental interaction: an interaction model for designing post-WIMP user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 446-453). ACM.

[B03] Baum, D. (2003). *Definitive Guide to Lego Mindstorms* (Vol. 2). Apress.

[B08] Barsalou, L. W. (2008). Grounded cognition. *Annu. Rev. Psychol.*, *59*, 617-645.

[B13] Berry, S. (2013, December 3). Connect 4. Retrieved July 25, 2014, from http://www.flickr.com/photos/unloveable/2396019222/.

[BB96] Bruns, W., & Brauer, V. (1996). *Bridging the Gap Between Real and Virtual Modeling: A New Approach to Human-computer Interaction*. Univ., Forschungszentrum Arbeit und Technik.

[BG10] Baskinger, M., & Gross, M. (2010). Tangible interaction= form+ computing.*interactions*, *17* (1), 6-11.

[BHA10] Bakker, S., van den Hoven, E., & Antle, A. N. (2011, January). MoSo tangibles: evaluating embodied learning. In *Proceedings of the fifth international conference on Tangible, embedded, and embodied interaction* (pp. 85-92). ACM.

[BJD04] Buur, J., Jensen, M. V., & Djajadiningrat, T. (2004, August). Hands-only scenarios and video action walls: novel methods for tangible user interaction design. In *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques* (pp. 185-192). ACM.

[BKP01] Billinghurst, M., Kato, H., & Poupyrev, I. (2001). The magicbook-moving seamlessly between reality and virtuality. *Computer Graphics and Applications, IEEE*, *21* (3), 6-8.

[BRS+03] Ballagas, R., Ringel, M., Stone, M., & Borchers, J. (2003, April). iStuff: a physical user interface toolkit for ubiquitous computing environments. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 537-544). ACM.

[C97] Clark, A. (1997). *Being there: Putting brain, body, and world together again*. MIT press.

[C01] Clark, A. (2001). Reasons, robots and the extended mind. *Mind & Language*, *16* (2), 121-145.

[C08] Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension: Embodiment, Action, and Cognitive Extension*. Oxford University Press.

[CC98] Clark, A., & Chalmers, D. (1998). The extended mind. *analysis*, 7-19.

[CDJ+02] Camarata, K., Do, E. Y. L., Johnson, B. R., & Gross, M. D. (2002, January). Navigational blocks: navigating information space with tangible media. In *Proceedings of the 7th international conference on Intelligent user interfaces* (pp. 31-38).

[CH93] Cypher, A., & Halbert, D. C. (Eds.). (1993). *Watch what I do: programming by demonstration*. MIT press.

[CMN86] Card, S. K., Moran, T. P., & Newell, A. (1986). The model human processor: An engineering model of human performance. *Handbook of Human Perception*, *2*.

[CRR08] Couture, N., Rivière, G., & Reuter, P. (2008, February). GeoTUI: a tangible user interface for geoscience. In *Proceedings of the 2nd international conference on Tangible and embedded interaction* (pp. 89-96).

[CWP99] Cohen, J., Withgott, M., & Piernot, P. (1999, May). Logjam: a tangible multi-person interface for video logging. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (pp. 128-135). ACM.

[D01] Paul Dourish. (2004). *Where the action is: the foundations of embodied interaction*. MIT press.

[DBS09] Doering, T., Beckhaus, S., & Schmidt, A. (2009, April). Towards a sensible integration of paper-based tangible user interfaces into creative work processes. In *CHI'09 Extended Abstracts on Human Factors in Computing Systems* (pp. 4627-4632).

[DPM+13] Donahue, T. J., Poor, G. M., Mott, M. E., Leventhal, L. M., Zimmerman, G., & Klopfer, D. (2013, February). On interface closeness and problem solving. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction* (pp. 139-146). ACM.

[DRL+11] Dijk, J. V., Roest, J. V. D., Lugt, R. V. D., & Overbeeke, K. C. (2011, November). NOOT: a tool for sharing moments of reflection during creative meetings. In *Proceedings of the 8th ACM conference on Creativity and cognition* (pp. 157-164). ACM.

[DWF+04] Djajadiningrat, T., Wensveen, S., Frens, J., & Overbeeke, K. (2004). Tangible products: redressing the balance between appearance and action. *Personal and Ubiquitous Computing*, *8* (5), 294-30

[EB09] Edge, D., & Blackwell, A. F. (2009, February). Peripheral tangible interaction by analytic design. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction* (pp. 69-76). ACM.

[EFH+76] Ekstrom, R. B., French, J. W., Harman, H. H., & Dermen, D. (1976). *Manual for kit of factor referenced cognitive tests* (pp. 109-113). Educational Testing Service.

[EO10] Esteves, A., & Oakley, I. (2010, October). Mementos: a tangible interface supporting travel. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries* (pp. 643-646). ACM.

[F95] Frazer, J. (1995). An evolutionary architecture.

[F96] Fitzmaurice, G. W. (1996). *Graspable user interfaces* (Doctoral dissertation, University of Toronto).

[FB97] Fitzmaurice, G. W., & Buxton, W. (1997, March). An empirical evaluation of graspable user interfaces: towards specialized, space-multiplexed input. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems* (pp. 43-50). ACM.

[FFF80] Frazer, J. H., Frazer, J. M., & Frazer, P. A. (1980). Intelligent physical three-dimensional modelling system. *Computer Graphics*, *80*, 359-70.

[FGH+00] Fishkin, K. P., Gujar, A., Harrison, B. L., Moran, T. P., & Want, R. (2000). Embodied user interfaces for really direct manipulation. *Communications of the ACM*, *43*(9), 74-80.

[FIB95] Fitzmaurice, G. W., Ishii, H., & Buxton, W. A. (1995, May). Bricks: laying the foundations for graspable user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 442-449). ACM Press/Addison-Wesley Publishing Company.

[FKY05] Feijs, L., Kyffin, S., & Young, B. (2005). *Proceedings of Design and Semantics of Form and Movement* – DesForM 2005. Foreword. Koninklijke Philips Electronics N.V. Eindhoven. 3, 2005.

[FMI00] Frei, P., Su, V., Mikhak, B., & Ishii, H. (2000, April). Curlybot: designing a new class of computational toys. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 129-136). ACM.

[FT06] Fernaeus, Y., & Tholander, J. (2006, April). Finding design qualities in a tangible programming space. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (pp. 447-456). ACM.

[FTJ08] Fernaeus, Y., Tholander, J., & Jonsson, M. (2008). Beyond representations: towards an action-centric perspective on tangible interaction. *International Journal of Arts and Technology*, *1* (3), 249-267.

[FTJ08b] Fernaeus, Y., Tholander, J., & Jonsson, M. (2008, February). Towards a new set of ideals: consequences of the practice turn in tangible interaction. In *Proceedings of the 2nd international conference on Tangible and embedded interaction* (pp. 223-230).

[G91] Gaver, W. W. (1991, March). Technology affordances. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 79-84). ACM .

[G97] Glenberg, A. M. (1997). What memory is for: Creating meaning in the service of action. *Behavioral and brain sciences*, *20* (01), 41-50.

[G05] Goldin-Meadow, S. (2005). *Hearing gesture: How our hands help us think*. Harvard University Press.

[GBB+04] Gaver, W. W., Bowers, J., Boucher, A., Gellerson, H., Pennington, S., Schmidt, A., ... & Walker, B. (2004, April). The drift table: designing for ludic engagement. In *CHI'04 extended abstracts on Human factors in computing systems* (pp. 885-900). ACM.

[GF01] Greenberg, S., & Fitchett, C. (2001, November). Phidgets: easy development of physical interfaces through physical widgets. In *Proceedings of the 14th annual ACM symposium on User interface software and technology* (pp. 209-218). ACM.

[GSH+07] Girouard, A., Solovey, E. T., Hirshfield, L. M., Ecott, S., Shaer, O., & Jacob, R. J. (2007, February). Smart Blocks: a tangible mathematical manipulative. In *Proceedings of the 1st international conference on Tangible and embedded interaction* (pp. 183-186). ACM.

[GSO05] Gillet, A., Sanner, M., Stoffler, D., & Olson, A. (2005). Tangible augmented interfaces for structural molecular biology. *Computer Graphics and Applications, IEEE*, *25*(2), 13-17.

[H81] Hyde, J. S. (1981). How large are cognitive gender differences? A meta-analysis using! w^2 and d. *American Psychologist*, *36* (8), 892.

[H12] Hornecker, E. (2012, February). Beyond affordance: tangibles' hybrid nature. In *Proceedings of the Sixth International Conference on Tangible, Embedded and Embodied Interaction* (pp. 175-182). ACM.

[HAM+07] Hartmann, B., Abdulla, L., Mittal, M., & Klemmer, S. R. (2007, April). Authoring sensor-based interactions by demonstration with direct manipulation and

pattern recognition. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 145-154). ACM.

[HB06] Hornecker, E., & Buur, J. (2006, April). Getting a grip on tangible interaction: a framework on physical space and social interaction. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (pp. 437-446). ACM.

[HFA+07] van den Hoven, E., Frens, J., Aliakseyeu, D., Martens, J. B., Overbeeke, K., & Peters, P. (2007, February). Design research & tangible interaction. In *Proceedings of the 1st international conference on Tangible and embedded interaction* (pp. 109-115). ACM.

[HHO08] Hengeveld, B., Hummels, C., Overbeeke, K., Voort, R., van Balkom, H., & de Moor, J. (2008, February). Let me actuate you. In *Proceedings of the 2nd international conference on Tangible and embedded interaction* (pp. 159-166). ACM.

[HHO09] Hengeveld, B., Hummels, C., Overbeeke, K., Voort, R., van Balkom, H., & de Moor, J. (2009, February). Tangibles for toddlers learning language. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction* (pp. 161-168).

[HIW08] Hurtienne, J., Israel, J. H., & Weber, K. (2008, February). Cooking up real world business applications combining physicality, digitality, and image schemas. In *Proceedings of the 2nd international conference on Tangible and embedded interaction* (pp. 239-246). ACM.

[HJ07] Horn, M. S., & Jacob, R. J. (2007, April). Tangible programming in the classroom with tern. In *CHI'07 extended abstracts on Human factors in computing systems* (pp. 1965-1970). ACM.

[HL07] Hornbæk, K., & Law, E. L. C. (2007, April). Meta-analysis of correlations among usability measures. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 617-626). ACM.

[HM11] van den Hoven, E., & Mazalek, A. (2011). Grasping gestures: Gesturing with physical artifacts. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, *25*(03), 255-271.

[HPG+94] Hinckley, K., Pausch, R., Goble, J. C., & Kassell, N. F. (1994, April). Passive real-world interface props for neurosurgical visualization. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 452-458). ACM.

[HRL99] Holmquist, L. E., Redström, J., & Ljungstrand, P. (1999, January). Token-based access to digital information. In *Handheld and Ubiquitous Computing* (pp. 234-245). Springer Berlin Heidelberg.

[HS88] Hart, S. G., & Staveland, L. E. (1988). Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. *Advances in psychology*,*52*, 139-183.

[HSC+09] Horn, M. S., Solovey, E. T., Crouser, R. J., & Jacob, R. J. (2009, April). Comparing the use of tangible and graphical programming languages for informal science education. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 975-984). ACM.

[HSJ08] Horn, M. S., Solovey, E. T., & Jacob, R. J. (2008, June). Tangible programming and informal science learning: making TUIs work for museums. In *Proceedings of the 7th international conference on Interaction design and children* (pp. 194-201). ACM.

[HYG03] Huang, C. J., Yi-Luen Do, E., & Gross, D. (2003). MouseHaus Table: a Physical Interface for Urban Design. In *16th Annual ACM Symposium on User Interface Software and Technology (UIST)*.

[I08] Ishii, H. (2008, February). Tangible bits: beyond pixels. In *Proceedings of the 2nd international conference on Tangible and embedded interaction* (pp. xv-xxv). ACM.

[I08b] Ishii, H. (2008). The tangible user interface and its evolution. *Communications of the ACM*, *51* (6), 32-36.

[IU97] Ishii, H., & Ullmer, B. (1997, March). Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems* (pp. 234-241). ACM.

[J08] Jordà, S. (2008). On stage: the reactable and other musical tangibles go real.*International Journal of Arts and Technology*, *1* (3), 268-287.

[JGA+07] Jordà, S., Geiger, G., Alonso, M., & Kaltenbrunner, M. (2007, February). The reacTable: exploring the synergy between live music performance and tabletop tangible interfaces. In *Proceedings of the 1st international conference on Tangible and embedded interaction* (pp. 139-146). ACM.

[JGH+08] Jacob, R. J., Girouard, A., Hirshfield, L. M., Horn, M. S., Shaer, O., Solovey, E. T., & Zigelbaum, J. (2008, April). Reality-based interaction: a framework for post-WIMP interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 201-210). ACM.

[JIP+02] Jacob, R. J., Ishii, H., Pangaro, G., & Patten, J. (2002, April). A tangible interface for organizing information using a grid. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 339-346). ACM.

[K95] Kirsh, D. (1995, January). Complementary strategies: Why we use our hands when we think. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (pp. 212-217). Mahwah, NJ: Lawrence Erlbaum Associates.

[K95b] Kirsh, D. (1995). The intelligent use of space. *Artificial intelligence*, *73* (1), 31-68.

[K96] Kirsh, D. (1996). Adapting the environment instead of oneself.

[K97] Kirsh, D. (1997). Interactivity and multimedia interfaces. *Instructional Science*, *25* (2), 79-96.

[K99] Kirsh, D. (1999). Distributed cognition, coordination and environment design. In *Proceedings of the European conference on Cognitive Science* (pp. 1-11).

[K09] Kaltenbrunner, M. Website on Tangible Music. Read April 2009. http://modin.yuri.at/tangibles/.

[K09b] Kirsh, D. (2009). Projection, problem space and anchoring.

[K09c] Kirsh, D. (2009). Problem solving and situated cognition.

[KB99] Kato, H., & Billinghurst, M. (1999). Marker tracking and hmd calibration for a video-based augmented reality conferencing system. In *Augmented Reality, 1999.(IWAR'99) Proceedings. 2nd IEEE and ACM International Workshop on* (pp. 85-94).

[KBP+00] Kato, H., Billinghurst, M., Poupyrev, I., Imamoto, K., & Tachibana, K. (2000). Virtual object manipulation on a table-top AR environment. In *Augmented Reality, 2000. (ISAR 2000). Proceedings. IEEE and ACM International Symposium on* (pp. 111-119). IEEE.

[KHT06] Klemmer, S. R., Hartmann, B., & Takayama, L. (2006, June). How bodies matter: five themes for interaction design. In *Proceedings of the 6th conference on Designing Interactive systems* (pp. 140-149). ACM.

[KLL+04] Klemmer, S. R., Li, J., Lin, J., & Landay, J. A. (2004, April). Papier-Mache: toolkit support for tangible input. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 399-406). ACM.

[KM94] Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive science*, *18* (4), 513-549.

[KST+09] Kirk, D., Sellen, A., Taylor, S., Villar, N., & Izadi, S. (2009, September). Putting the physical into the digital: issues in designing hybrid interactive surfaces. In *Proceedings of the 23rd British HCI Group Annual Conference on People and Computers: Celebrating People and Technology* (pp. 35-44). British Computer Society.

[L93] Laurel, B. (2013). *Computers as theatre*. Pearson Education.

[LBB+07] van Loenen, E., Bergman, T., Buil, V., van Gelder, K., Groten, M., Hollemans, G., ... & van de Wijdeven, S. (2007, January). Entertaible: a solution for social gaming experiences. In *Tangible Play workshop, IUI Conference*.

[LGE08] Lin, A., Gregor, S., & Ewing, M. (2008). Developing a scale to measure the enjoyment of Web experiences. *Journal of Interactive Marketing*, *22*(4), 40-57.

[LHY+08] Leitner, J., Haller, M., Yun, K., Woo, W., Sugimoto, M., & Inami, M. (2008, December). IncreTable, a mixed reality tabletop game experience. In *Proceedings of the 2008 International Conference on Advances in Computer Entertainment Technology* (pp. 9-16). ACM.

[LJ80] Lakoff, G., & Johnson, M. (1980). The metaphorical structure of the human conceptual system. *Cognitive Science*, *4*(2), 195-208.

[LJ80b] Lakoff, G., & Johnson, M. (2008). *Metaphors we live by*. University of Chicago press.

[LK77] Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *biometrics*, 159-174.

[LNB+04] Lee, G. A., Nelles, C., Billinghurst, M., & Kim, G. J. (2004, November). Immersive authoring of tangible augmented reality applications. In *Proceedings of the 3rd IEEE/ACM international Symposium on Mixed and Augmented Reality* (pp. 172-181). IEEE Computer Society.

[LPI07] LeClerc, V., Parkes, A., & Ishii, H. (2007, April). Senspectra: a computationally augmented physical modeling toolkit for sensing and visualization of structural strain. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 801-804). ACM.

[LPI11] Lee, J., Post, R., & Ishii, H. (2011, October). ZeroN: mid-air tangible interaction enabled by computer controlled magnetic levitation. In *Proceedings of the 24th annual ACM symposium on User interface software and technology* (pp. 327-336).

[M92] MacKenzie, I. S. (1992). Fitts' law as a research and design tool in human-computer interaction. *Human-computer interaction*, *7* (1), 91-139.

[M99] MacKay, W. E. (1999). Is paper safer? The role of paper flight strips in air traffic control. *ACM Transactions on Computer-Human Interaction (TOCHI)*,*6*(4), 311-340.

[M00] Martin, F. G. (2000). *Robotic explorations: a hands-on introduction to engineering*. Prentice Hall PTR.

[M04] McNerney, T. S. (2004). From turtles to Tangible Programming Bricks: explorations in physical language design. *Personal and Ubiquitous Computing*, *8* (5), 326-337.

[M06] Miller, B. J. (2006). Cognition and the City.

[M07] Marshall, P. (2007, February). Do tangible interfaces enhance learning?. In *Proceedings of the 1st international conference on Tangible and embedded interaction* (pp. 163-170). ACM.

[MBS92] Mills, D., Bonn, E. A., & San Juan, S. S. T. M. (1992). Macintosh Human Interface Guidelines.

[MCL10] Marshall, P., Cheng, P. C. H., & Luckin, R. (2010). Tangibles in the balance: a discovery learning task with physical or graphical materials. In *Proceedings of the fourth international conference on Tangible, embedded, and embodied interaction* (pp. 153-160). ACM.

[MF99] Mackay, W. E., & Fayard, A. L. (1999, November). Designing interactive paper: lessons from three augmented reality projects. In *Proceedings of the international workshop on Augmented reality: placing artificial objects in real scenes: placing artificial objects in real scenes* (pp. 81-90). AK Peters, Ltd..

[MH09] Mazalek, A., & Van Den Hoven, E. (2009). Framing tangible interaction frameworks. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, *23* (03), 225-235.

[MKM07] Merrill, D., Kalanithi, J., & Maes, P. (2007, February). Siftables: towards sensor network user interfaces. In *Proceedings of the 1st international conference on Tangible and embedded interaction* (pp. 75-78). ACM.

[MMS00] Martin, F., Mikhak, B., & Silverman, B. (2000). MetaCricket: A designer's kit for making computational devices. *IBM Systems Journal*, *39* (3.4), 795-815.

[N88] Norman, D. A. (1988). *The psychology of everyday things*. Basic books.

[N10] Norman, D. A. (2010). Natural user interfaces are not natural. *interactions*, *17* (3), 6-10.

[NK14] Nowacka, D., & Kirk, D. (2014, February). Tangible autonomous interfaces (tais): Exploring autonomous behaviours in tuis. In *Proceedings of the 8th International Conference on Tangible, Embedded and Embodied Interaction* (pp. 1-8). ACM.

[NL95] Newman, W. M., Lamming, M. G., & Lamming, M. (1995). *Interactive system design* (pp. I-XXXIV). Reading: Addison-Wesley.

[NNG03] Newton-Dunn, H., Nakano, H., & Gibson, J. (2003, May). Block jam: a tangible interface for interactive music. In *Proceedings of the 2003 conference on New interfaces for musical expression* (pp. 170-177). National University of Singapore.

[O92] Olsen, D. (1992). *User interface management systems: models and algorithms*. Morgan Kaufmann Publishers Inc..

[O99] Oviatt, S. (1999). Ten myths of multimodal interaction. *Communications of the ACM*, *42* (11), 74-81.

[O13] Osborn, S. (2013). *Makers at Work: Folks Reinventing the World One Object Or Idea at a Time*. Apress.

[OF04] O'Malley, C., & Fraser, D. S. (2004). Literature review in learning with tangible technologies.

[P95] Petre, M. (1995). Why looking isn't always seeing: readership skills and graphical programming. *Communications of the ACM*, *38* (6), 33-44.

[P99] Paradiso, J. A. (1999). The brain opera technology: New instruments and gestural sensors for musical interaction and performance. *Journal of New Music Research*, *28* (2), 130-149.

[PFD+95] Parsons, L. M., Fox, P. T., Downs, J. H., Glass, T., Hirsch, T. B., Martin, C. C., Jerabek, Pl A., & Lancaster, J. L. (1995). *Use of implicit motor imagery for visual shape discrimination as revealed by PET*. Nature, 375, 54-58.

[PH09] Pedersen, E. W., & Hornbæk, K. (2009, February). mixiTUI: a tangible sequencer for electronic live performances. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction* (pp. 223-230). ACM.

[PI00] Patten, J., & Ishii, H. (2000, April). A comparison of spatial organization strategies in graphical and tangible user interfaces. In *Proceedings of DARE 2000 on Designing augmented reality environments* (pp. 41-50). ACM.

[PI07] Patten, J., & Ishii, H. (2007, April). Mechanical constraints as computational constraints in tabletop tangible interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 809-818). ACM.

[PNO07] Poupyrev, I., Nashida, T., & Okabe, M. (2007, February). Actuation and tangible user interfaces: the Vaucanson duck, robots, and shape displays. In *Proceedings of the 1st international conference on Tangible and embedded interaction* (pp. 205-212). ACM.

[PRI02] Patten, J., Recht, B., & Ishii, H. (2002, May). Audiopad: a tag-based interface for musical performance. In *Proceedings of the 2002 conference on New interfaces for musical expression* (pp. 1-6). National University of Singapore.

[PW92] Proctera, R. N., & Williamsb, R. A. (1992). HCI: whose problem is IT anyway?.

[R93] Resnick, M. (1993). Behavior construction kits. *Communications of the ACM*, *36* (7), 64-71.

[R06] Rogers, Y. (2006). Moving on from weiser's vision of calm computing: Engaging ubicomp experiences. In *UbiComp 2006: Ubiquitous Computing* (pp. 404-421). Springer Berlin Heidelberg.

[RGB+10] Raynal, M., Gauffre, G., Bach, C., Schmitt, B., & Dubois, E. (2010, October). Tactile camera vs. tangible camera: taking advantage of small physical artefacts to navigate into large data collection. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries* (pp. 373-382). ACM.

[RFK+98] Rauterberg, M., Fjeld, M., Krueger, H., Bichsel, M., Leonhardt, U., & Meier, M. (1998, April). BUILD-IT: a planning tool for construction and design. In *CHI 98 Cconference Summary on Human Factors in Computing Systems* (pp. 177-178). ACM.

[RMB+98] Resnick, M., Martin, F., Berg, R., Borovoy, R., Colella, V., Kramer, K., & Silverman, B. (1998, January). Digital manipulatives: new toys to think with. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 281-287). ACM Press/Addison-Wesley Publishing Co..

[RMI04] Ryokai, K., Marti, S., & Ishii, H. (2004, April). I/O brush: drawing with everyday objects as ink. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 303-310). ACM.

[RPI04] Raffle, H. S., Parkes, A. J., & Ishii, H. (2004, April). Topobo: a constructive assembly system with kinetic memory. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 647-654). ACM.

[S99] Seitz, J. A. (2000). The bodily basis of thought. *New ideas in Psychology*, *18* (1), 23-40.

[S99b] Schwartz, D. L. (1999). Physical imagery: Kinematic versus dynamic models. *Cognitive Psychology*, *38*(3), 433-464.

[S07] Smith, A. C. (2007, February). Using magnets in physical blocks that behave as programming objects. In *Proceedings of the 1st international conference on Tangible and embedded interaction* (pp. 147-150). ACM.

[S10] Shapiro, L. A. (2010). *Embodied Cognition (New Problems of Philosophy)*. London: Routledge.

[S14] Seeeduino Film. (n.d.). Retrieved November 25, 2014, from http://www.seeedstudio.com/depot/Seeeduino-Film-p-689.html

[SA07] Schacter, D. L., & Addis, D. R. (2007). The cognitive neuroscience of constructive memory: remembering the past and imagining the future. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362* (1481), 773-786.

[SBG+08] Sturm, J., Bekker, T., Groenendaal, B., Wesselink, R., & Eggen, B. (2008, June). Key issues for the successful design of an intelligent, interactive playground. In *Proceedings of the 7th international conference on Interaction design and children* (pp. 258-265). ACM.

[SCB07] Sitorus, L., Cao, S. S., & Buur, J. (2007, February). Tangible user interfaces for configuration practices. In *Proceedings of the 1st international conference on Tangible and embedded interaction* (pp. 223-230). ACM.

[SCI04] Scott, S. D., Carpendale, M. S. T., & Inkpen, K. M. (2004, November). Territoriality in collaborative tabletop workspaces. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work* (pp. 294-303). ACM.

[SDM+01] Sherman, L., Druin, A., Montemayor, J., Farber, A., Platner, M., Simms, S., Porteous, J., Alborzi, H., Best, J., Hammer, J., Kruskal, A., Matthews, J., Rhodes, E., Cosans, C. & Lal, A. (2001, March). StoryKit: tools for children to build room-sized interactive experiences. In *CHI'01 Extended Abstracts on Human Factors in Computing Systems* (pp. 197-198). ACM.

[SH10] Shaer, O., & Hornecker, E. (2010). Tangible user interfaces: past, present, and future directions. *Foundations and Trends in Human-Computer Interaction*, *3*(1–2), 1-137.

[SIW+02] Sharlin, E., Itoh, Y., Watson, B., Kitamura, Y., Sutphen, S., & Liu, L. (2002, April). Cognitive cubes: a tangible user interface for cognitive assessment. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 347-354). ACM.

[SJ09] Shaer, O., & Jacob, R. J. (2009). A specification paradigm for the design and implementation of tangible user interfaces. *ACM Transactions on Computer-Human Interaction (TOCHI)*, *16* (4), 20.

[SK93] Suzuki, H., & Kato, H. (1993, August). AlgoBlock: a tangible programming language, a tool for collaborative learning. In *Proceedings of 4th European Logo Conference* (pp. 297-303).

[SK95] Suzuki, H., & Kato, H. (1995, October). Interaction-level support for collaborative learning: AlgoBlock—an open programming language. In *The first international conference on Computer support for collaborative learning* (pp. 349-355). L. Erlbaum Associates Inc..

[SKM09] Soute, I., Kaptein, M., & Markopoulos, P. (2009, June). Evaluating outdoor play for children: virtual vs. tangible game objects in pervasive games. In *Proceedings of the 8th International Conference on Interaction Design and Children* (pp. 250-253).

[SLC+04] Shaer, O., Leland, N., Calvillo-Gamez, E. H., & Jacob, R. J. (2004). The TAC paradigm: specifying tangible user interfaces. *Personal and Ubiquitous Computing*, *8* (5), 359-369.

[SSV+11] Shaer, O., Strait, M., Valdes, C., Feng, T., Lintz, M., & Wang, H. (2011, May). Enhancing genomic learning through tabletop interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2817-2826). ACM.

[SV08] Schiettecatte, B., & Vanderdonckt, J. (2008, February). AudioCubes: a distributed cube tangible interface based on interaction range for sound design. In *Proceedings of the 2nd international conference on Tangible and embedded interaction* (pp. 3-10). ACM.

[T85] Tarjan, R. E. (1985). Amortized computational complexity. *SIAM Journal on Algebraic Discrete Methods*, *6* (2), 306-318.

[T07] Turkle, S. (Ed.). (2007). *Evocative objects: Things we think with*. MIT press.

[TKI10] Tuddenham, P., Kirk, D., & Izadi, S. (2010, April). Graspables revisited: multi-touch vs. tangible input for tabletop displays in acquisition and manipulation tasks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2223-2232). ACM.

[TT06] Toney, A., & Thomas, B. H. (2006, November). Applying reach in direct manipulation user interfaces. In *Proceedings of the 18th Australia conference on Computer-Human Interaction: Design: Activities, Artefacts and Environments* (pp. 393-396). ACM.

[UI98] Underkoffler, J., & Ishii, H. (1998, January). Illuminating light: An optical design tool with a luminous-tangible interface. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 542-549). ACM Press/Addison-Wesley Publishing Co..

[UI99] Underkoffler, J., & Ishii, H. (1999, May). Urp: a luminous-tangible workbench for urban planning and design. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (pp. 386-393). ACM.

[UI99b] Ullmer, B., & Ishii, H. (1999, May). mediaBlocks: tangible interfaces for online media. In *CHI'99 extended abstracts on Human factors in computing systems* (pp. 31-32). ACM.

[UI00] Ullmer, B., & Ishii, H. (2000). Emerging frameworks for tangible user interfaces.*IBM systems journal*, *39* (3.4), 915-931.

[UIJ05] Ullmer, B., Ishii, H., & Jacob, R. J. (2005). Token + constraint systems for tangible interaction with digital information. *ACM Transactions on Computer-Human Interaction (TOCHI)*, *12* (1), 81-118.

[VVB95] Voyer, D., Voyer, S., & Bryden, M. P. (1995). Magnitude of sex differences in spatial abilities: a meta-analysis and consideration of critical variables. *Psychological bulletin*, *117* (2), 250.

[VP05] Vallée-Tourangeau, F., & Krüsi Penney, A. (2005). The impact of external representation in a rule discovery task. *European Journal of Cognitive Psychology*, *17* (6), 820-834.

[VSK08] Virnes, M., Sutinen, E., & Kärnä-Lin, E. (2008, June). How children's individual needs challenge the design of educational robotics. In *Proceedings of the 7th international conference on Interaction design and children* (pp. 274-281). ACM.

[VWH05] Vallabhaneni, A., Wang, T., & He, B. (2005). Brain—computer interface. In *Neural Engineering* (pp. 85-121). Springer US.

[W93] Weiser, M. (1993). Hot topics-ubiquitous computing. *Computer*, *26* (10), 71-72.

[W93b] Weiser, M. (1993). Some computer science issues in ubiquitous computing. *Communications of the ACM*, *36* (7), 75-84.

[W02] Wilson, M. (2002). Six views of embodied cognition. *Psychonomic bulletin & review*, *9* (4), 625-636.

[WIH+08] Wilson, A. D., Izadi, S., Hilliges, O., Garcia-Mendoza, A., & Kirk, D. (2008, October). Bringing physics to the surface. In *Proceedings of the 21st annual ACM symposium on User interface software and technology* (pp. 67-76). ACM.

[WB96] Weiser, M., & Brown, J. S. (1996). Designing calm technology. *PowerGrid Journal*, *1* (1), 75-85.

[WDO04] Wensveen, S. A., Djajadiningrat, J. P., & Overbeeke, C. J. (2004, August). Interaction frogger: a design framework to couple action and function through

feedback and feedforward. In *Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques* (pp. 177-184). ACM.

[WG01] Weinberg, G., & Gan, S. L. (2001). The squeezables: Toward an expressive and interdependent multi-player musical instrument. *Computer Music Journal*, *25* (2), 37-45.

[WKS08] Westeyn, T. L., Kientz, J. A., Starner, T. E., & Abowd, G. D. (2008, June). Designing toys with automatic play characterization for supporting the assessment of a child's development. In *Proceedings of the 7th international conference on Interaction design and children* (pp. 89-92). ACM.

[WMG93] Wellner, P., Mackay, W., & Gold, R. (1993). Back to the real world. *Communications of the ACM*, *36* (7), 24-26.

[WP02] Wyeth, P., & Purchase, H. C. (2002, April). Tangible programming elements for young children. In *CHI'02 Extended Abstracts on Human Factors in Computing Systems* (pp. 774-775). ACM.

[WVV11] Weller, A., Villejoubert, G., & Vallée-Tourangeau, F. (2011). Interactive insight problem solving. *Thinking & Reasoning*, *17* (4), 424-439.

[XAM08] Xie, L., Antle, A. N., & Motamedi, N. (2008, February). Are tangibles more fun?: comparing children's enjoyment and engagement using physical, graphical and tangible user interfaces. In *Proceedings of the 2nd international conference on Tangible and embedded interaction* (pp. 191-198). ACM.

[Z97] Zhang, J. (1997). The nature of external representations in problem solving. *Cognitive science*, *21* (2), 179-217.

[Z12] Zhang, Z. (2012). Microsoft kinect sensor and its effect. *MultiMedia, IEEE*, *19* (2), 4-10.

[ZAM+12] Zaman, B., Vanden Abeele, V., Markopoulos, P., & Marshall, P. (2012). Editorial: the evolving field of tangible interaction for children: the challenge of empirical validation. *Personal and Ubiquitous Computing*, *16* (4), 367-378.

[ZAR05] Zuckerman, O., Arida, S., & Resnick, M. (2005, April). Extending tangible interfaces for education: digital montessori-inspired manipulatives. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 859-868). ACM.

[ZJL+09] Zufferey, G., Jermann, P., Lucchi, A., & Dillenbourg, P. (2009, February). TinkerSheets: using paper forms to control and visualize tangible simulations. In *Proceedings of the 3rd international Conference on Tangible and Embedded interaction* (pp. 377-384). ACM.

[ZN94] Zhang, J., & Norman, D. A. (1994). Representations in distributed cognitive tasks. *Cognitive science*, *18* (1), 87-122.

ANNEXES

Annex 1

The entity–relationship diagram of tiREC (chapter four)

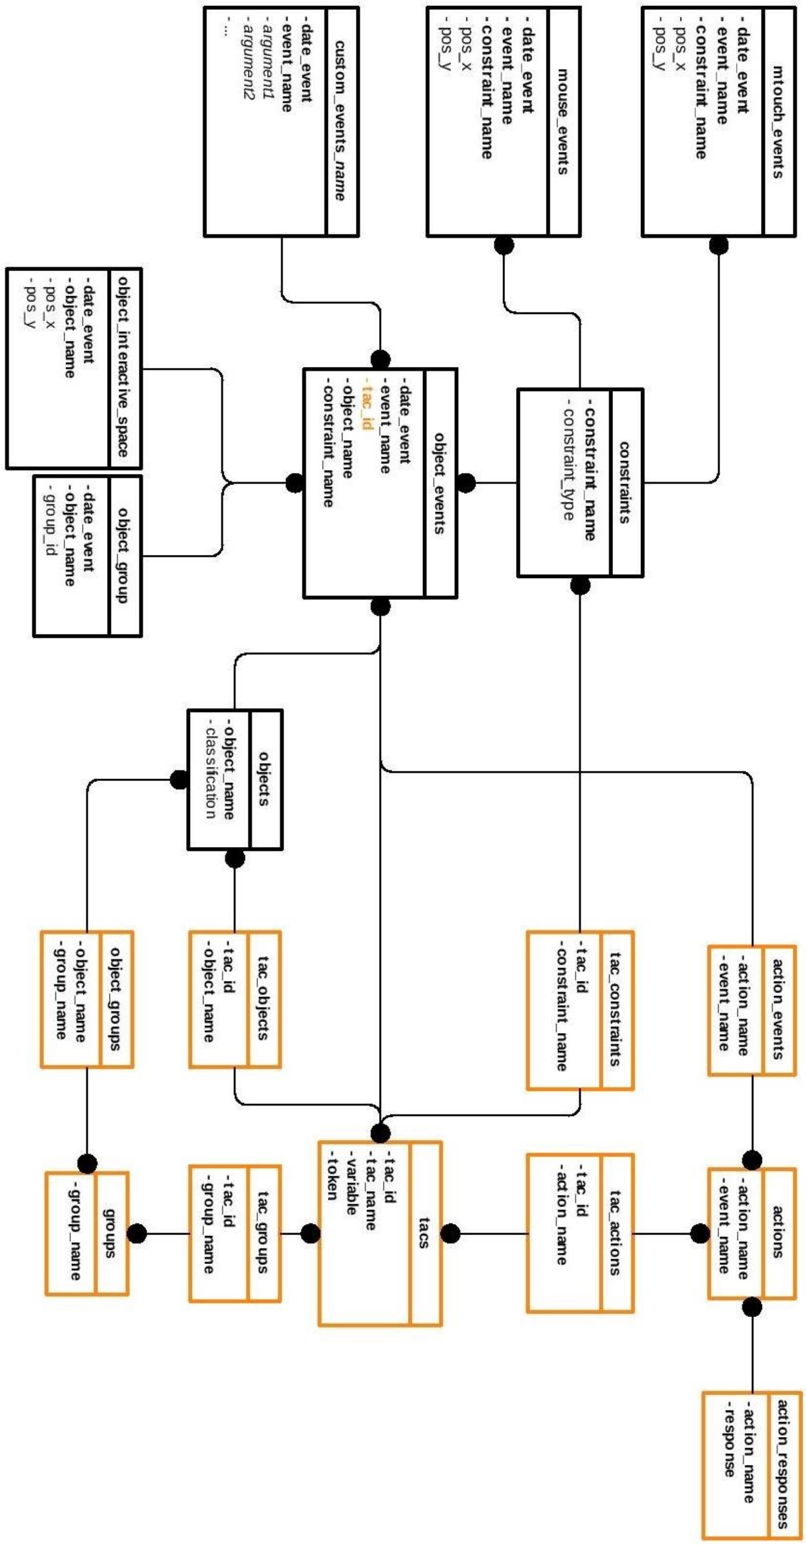

Annex 2

The list of references that informed the ATB framework (chapter seven)

Abrams, R. (1999). Adventures in tangible computing: The work of interaction designer 'Durrell Bishop'in context. magistrska naloga, Royal College of Art, London.

Alibali, M. W., & DiRusso, A. A. (1999). The function of gesture in learning to count: More than keeping track. Cognitive development, 14(1), 37-56.

Antle, A. N. (2007, February). The CTI framework: informing the design of tangible systems for children. In Proceedings of the 1st international conference on Tangible and embedded interaction (pp. 195-202). ACM.

Antle, A. N. (2013). Research opportunities: Embodied child–computer interaction. International Journal of Child-Computer Interaction, 1(1), 30-36.

Antle, A. N. (2013). Exploring how children use their hands to think: An embodied interactional analysis. Behaviour & Information Technology, 32(9), 938-954.

Antle, A. N., Corness, G., & Bevans, A. (2013). Balancing justice: comparing whole body and controller–based interaction for an abstract domain. International Journal of Arts and Technology, 6(4), 388-409.

Antle, A. N., Droumeva, M., & Ha, D. (2009, April). Thinking with hands: an embodied approach to the analysis of children's interaction with computational objects. In CHI'09 Extended Abstracts on Human Factors in Computing Systems (pp. 4027-4032). ACM.

Antle, A. N., Droumeva, M., & Ha, D. (2009, June). Hands on what?: comparing children's mouse-based and tangible-based interaction. In Proceedings of the 8th International Conference on Interaction Design and Children (pp. 80-88). ACM.

Antle, A. N., & Wang, S. (2013, February). Comparing motor-cognitive strategies for spatial problem solving with tangible and multi-touch interfaces. In Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction (pp. 65-72). ACM.

Balakrishnan, R., & Kurtenbach, G. (1999, May). Exploring bimanual camera control and object manipulation in 3D graphics interfaces. In Proceedings of the SIGCHI conference on Human Factors in Computing Systems (pp. 56-62). ACM.

Ballard, D., Hayhoe, M., & Pelz, J. (1995). Memory representations in natural tasks. Cognitive Neuroscience, Journal of, 7(1), 66-80.

Carlson, R. A., Avraamides, M. N., Cary, M., & Strasberg, S. (2007). What do the hands externalize in simple arithmetic?. Journal of Experimental Psychology: Learning, Memory, and Cognition, 33(4), 747.

Cary, M., & Carlson, R. A. (2001). Distributing working memory resources during problem solving. Journal of Experimental Psychology: Learning, Memory, and Cognition, 27(3), 836.

Clark, A. (1997). Being there: Putting brain, body, and world together again. MIT press.

Clark, A., & Chalmers, D. (1998). The extended mind. analysis, 7-19.

Cohen, J., Withgott, M., & Piernot, P. (1999, May). Logjam: a tangible multi-person interface for video logging. In Proceedings of the SIGCHI conference on Human Factors in Computing Systems (pp. 128-135). ACM.

Collins, A., & Ferguson, W. (1993). Epistemic forms and epistemic games: Structures and strategies to guide inquiry. Educational Psychologist, 28(1), 25-42.

Cox, R. (1999). Representation construction, externalised cognition and individual differences. Learning and instruction, 9(4), 343-363.

De Cruz, H., & De Smedt, J. (2013). Mathematical symbols as epistemic actions. Synthese, 190(1), 3-19.

Destefano, M., Lindstedt, J., & Gray, W. (2011). Use of complementary actions decreases with expertise. In Proceedings of the 33rd Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society (pp. 2709-2014).

Donahue, T. J., Poor, G. M., Mott, M. E., Leventhal, L. M., Zimmerman, G., & Klopfer, D. (2013, February). On interface closeness and problem solving. In Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction (pp. 139-146). ACM.

Esteves, A., & Oakley, I. (2010, October). Mementos: a tangible interface supporting travel. In Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries (pp. 643-646). ACM.

Esteves, A., Scott, M., & Oakley, I. (2013, February). Supporting offline activities on interactive surfaces. In Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction (pp. 147-154). ACM.

Esteves, A., van den Hoven, E., & Oakley, I. (2013, February). Physical games or digital games?: Comparing support for mental projection in tangible and virtual representations of a problem-solving task. In Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction (pp. 167-174). ACM.

Funk, M., Brugger, P., & Wilkening, F. (2005). Motor processes in children's imagery: The case of mental rotation of hands. Developmental Science, 8(5), 402-408.

Glenberg, A. M., & Robertson, D. A. (1999). Indexical understanding of instructions. Discourse Processes, 28(1), 1-26.

Goldin-Meadow, S. (2005). Hearing gesture: How our hands help us think. Harvard University Press.

Goodyear, P., Ellis, R., Brew, A., & Sachs, J. (2007). The development of epistemic fluency: Learning to think for a living.

Greeno, J. G., Riley, M. S., & Gelman, R. (1984). Conceptual competence and children's counting. Cognitive Psychology, 16(1), 94-143.

Hershkowitz, R., Schwarz, B. B., & Dreyfus, T. (2001). Abstraction in context: Epistemic actions. Journal for Research in Mathematics Education, 195-222.

Hitch, G. J. (1978). The role of short-term working memory in mental arithmetic. Cognitive Psychology, 10(3), 302-323.

Hornecker, E., & Buur, J. (2006, April). Getting a grip on tangible interaction: a framework on physical space and social interaction. In Proceedings of the SIGCHI conference on Human Factors in computing systems (pp. 437-446). ACM.

Hutchins, E. (1995). Cognition in the Wild (Vol. 262082314). Cambridge, MA: MIT press.

Hutchins, E. (1995). How a cockpit remembers its speeds. Cognitive science, 19(3), 265-288.

Jensen, M. V., & Stienstra, M. (2007, August). Making sense: Interactive sculptures as tangible design material. In Proceedings of the 2007 conference on Designing pleasurable products and interfaces (pp. 255-269). ACM.

Kastens, K. A., Liben, L. S., & Agrawal, S. (2008). Epistemic actions in science education. In Spatial Cognition VI. Learning, Reasoning, and Talking about Space (pp. 202-215). Springer Berlin Heidelberg.

Kirlik, A. (1998, March). The ecological expert: Acting to create information to guide action. In Human Interaction with Complex Systems, 1998. Proceedings., Fourth Annual Symposium on (pp. 15-27). IEEE.

Kirsh, D. (1995). The intelligent use of space. Artificial intelligence, 73(1), 31-68.

Kirsh, D. (1995, January). Complementary strategies: Why we use our hands when we think. In Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society (pp. 212-217). Mahwah, NJ: Lawrence Erlbaum Associates.

Kirsh, D. (1996). Adapting the environment instead of oneself.

Kirsh, D. (1997). Interactivity and multimedia interfaces. Instructional Science, 25(2), 79-96.

Kirsh, D. (1999). Distributed cognition, coordination and environment design. In Proceedings of the European conference on Cognitive Science (pp. 1-11).

Kirsh, D. (2000). A few thoughts on cognitive overload.

Kirsh, D. (2006). Distributed cognition: A methodological note. Pragmatics & Cognition, 14(2), 249-262.

Kirsh, D. (2010). Comparing tangible and virtual exploration of archaeological objects.

Kirsh, D. (2011). Situating Instructions.

Kirsh, D. (2013). Thinking with external representations. In Cognition Beyond the Brain (pp. 171-194). Springer London.

Kirsh, D., & Maglio, P. (1992). Some Epistemic Benefits of Action-Tetris, a Case Study.

Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. Cognitive science, 18(4), 513-549.

Klemmer, S. R., Hartmann, B., & Takayama, L. (2006, June). How bodies matter: five themes for interaction design. In Proceedings of the 6th conference on Designing Interactive systems (pp. 140-149). ACM.

Liben, L. S., Myers, L. J., & Kastens, K. A. (2008). Locating oneself on a map in relation to person qualities and map characteristics. In Spatial Cognition VI. Learning, Reasoning, and Talking about Space (pp. 171-187). Springer Berlin Heidelberg.

Maglio, P. P., & Kirsh, D. (1996). Epistemic action increases with skill. In Proceedings of the eighteenth annual conference of the cognitive science society (Vol. 16, pp. 391-396).

Maglio, P. P., & Wenger, M. J. (2000, August). Two views are better than one: Epistemic actions may prime. In Proceedings of the Twenty-second Annual Conference of the Cognitive Science Society, Mahwah, NJ. LEA.

Maglio, P. P., Wenger, M. J., & Copeland, A. M. (2003). The benefits of epistemic action outweigh the costs. In Proceedings of the 25th annual conference of the Cognitive Science Society (pp. 752-757).

Maglio, P. P., Wenger, M. J., & Copeland, A. M. (2008). Evidence for the role of self-priming in epistemic action: Expertise and the effective use of memory. Acta Psychologica, 127(1), 72-88.

Manches, A., O'Malley, C., & Benford, S. (2010). The role of physical representations in solving number problems: A comparison of young children's use of physical and virtual materials. Computers & Education, 54(3), 622-640.

Menary, R. (2010). Dimensions of mind. Phenomenology and the Cognitive Sciences, 9(4), 561-578.

Naveh-Benjamin, M. (1987). Coding of spatial location information: An automatic process?. Journal of Experimental Psychology: Learning, Memory, and Cognition, 13(4), 595.

Neth, H., & Payne, S. J. (2001). Addition as interactive problem solving. In Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society (pp. 698-703).

Neth, H., & Payne, S. J. (2011). Interactive coin addition: How hands can help us think. Proceedings of the 33rd CogSci, 279-284.

Nurminen, A., & Oulasvirta, A. (2008). Designing interactions for navigation in 3D mobile maps. In Map-based mobile services (pp. 198-227). Springer Berlin Heidelberg.

Owen, R., Kurtenbach, G., Fitzmaurice, G., Baudel, T., & Buxton, B. (2005, May). When it gets more difficult, use both hands: exploring bimanual curve manipulation. In Proceedings of Graphics Interface 2005 (pp. 17-24). Canadian Human-Computer Communications Society.

Piaget, J. (1980). Adaptation and intelligence (S. Eames, trans.).

Petre, M. (1995). Why looking isn't always seeing: readership skills and graphical programming. Communications of the ACM, 38(6), 33-44.

Resnick, M. (2006). Computer as paint brush: Technology, play, and the creative society. Play= learning: How play motivates and enhances children's cognitive and social-emotional growth, 192-208.

Sellen, A. J., & Harper, R. H. (2002). The myth of the paperless office. MIT press.

Sève, C., Saury, J., Leblanc, S., & Durand, M. (2005). Course-of-action theory in table tennis: a qualitative analysis of the knowledge used by three elite players during matches. Revue Européenne de Psychologie Appliquée/European Review of Applied Psychology, 55(3), 145-155.

Siegler, R. S. (1987). Strategy choices in subtraction.

Skundergard, K., Kirsh, D., & Dahlbäck, N. (2012). Maps in the Head and Maps in the Hand. In CogSci 2012 (pp. 2339-2344).

Tarjan, R. E. (1985). Amortized computational complexity. SIAM Journal on Algebraic Discrete Methods, 6(2), 306-318.

Terrenghi, L., Kirk, D., Richter, H., Krämer, S., Hilliges, O., & Butz, A. (2008, May). Physical handles at the interactive surface: exploring tangibility and its benefits. In Proceedings of the working conference on Advanced visual interfaces (pp. 138-145). ACM.

Trudel, C. I., & Payne, S. J. (1995). Reflection and goal management in exploratory learning. International Journal of Human-Computer Studies, 42(3), 307-339.

Uttal, D. H., Scudder, K. V., & DeLoache, J. S. (1997). Manipulatives as symbols: A new perspective on the use of concrete objects to teach mathematics. Journal of applied developmental psychology, 18(1), 37-54.

Vygotsky, L. S. (1980). Mind in society: The development of higher psychological processes. Harvard university press.

Wakkary, R., & Hatala, M. (2007). Situated play in a tangible interface and adaptive audio museum guide. Personal and Ubiquitous Computing, 11(3), 171-191.

Whitehead, A. N. (1911). An introduction to mathematics.

Wilson, M. (2002). Six views of embodied cognition. Psychonomic bulletin & review, 9(4), 625-636.

Wright, P. C., Fields, R. E., & Harrison, M. D. (2000). Analyzing human-computer interaction as distributed cognition: the resources model. Human-Computer Interaction, 15(1), 1-41.

Annex 3

The ATB framework scheme (chapter seven)

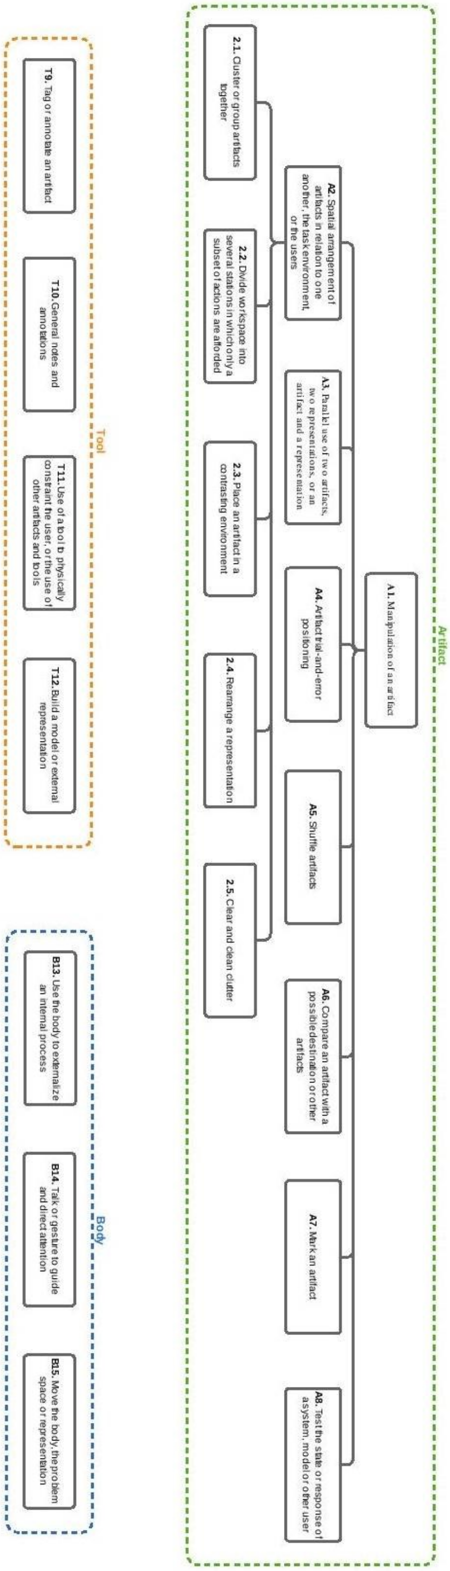

Annex 4

How to video-code using the ATB framework (chapter seven)

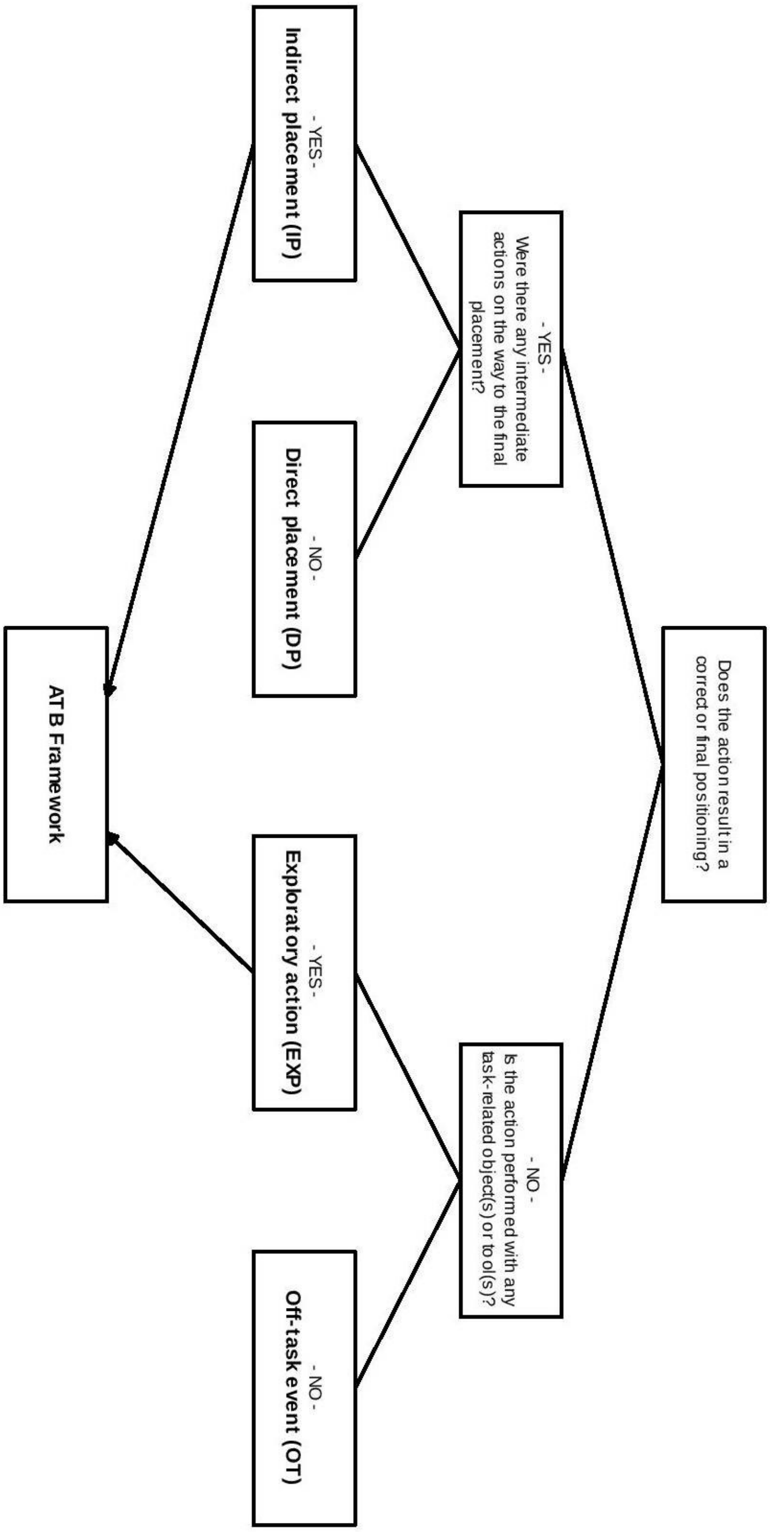

Annex 5

The ATB framework coding scheme file for Anvil (chapter seven)

```xml
<?xml version="1.0" encoding="ISO-8859-1"?>
<annotation-spec>
  <body>
        <group name="Track1">
                <track-spec name="Action" type="primary" color-attr="action">
                        <attribute name="action">
                                <value-el color="yellow">DP</value-el>
                                <value-el color="green">IP</value-el>
                                <value-el color="blue">Exploratory</value-el>
                                <value-el color="orange">Off-task</value-el>
                        </attribute>
                </track-spec>
                <track-spec name="EAs" type="span" ref="Track1.Action">
                        <attribute name="EAs" valuetype="String" />
                </track-spec>
        </group>
        <group name="Track2">
                <track-spec name="Action" type="primary" color-attr="action">
                        <attribute name="action">
                                <value-el color="yellow">DP</value-el>
                                <value-el color="green">IP</value-el>
                                <value-el color="blue">Exploratory</value-el>
                                <value-el color="orange">Off-task</value-el>
                        </attribute>
                </track-spec>
                <track-spec name="EAs" type="span" ref="Track2.Action">
                        <attribute name="EAs" valuetype="String" />
                </track-spec>
        </group>
        <group name="Track3">
                <track-spec name="Action" type="primary" color-attr="action">
                        <attribute name="action">
                                <value-el color="yellow">DP</value-el>
                                <value-el color="green">IP</value-el>
                                <value-el color="blue">Exploratory</value-el>
                                <value-el color="orange">Off-task</value-el>
                        </attribute>
                </track-spec>
                <track-spec name="EAs" type="span" ref="Track3.Action">
                        <attribute name="EAs" valuetype="String" />
                </track-spec>
        </group>
  </body>
</annotation-spec>
```