University of Alberta

Blurring the Boundaries: Spatial Tangible User Interfaces for Cognitive Assessment

by

Ehud Sharlin

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Department of Computing Science

Edmonton, Alberta
Fall 2003

Your file Votre référence
ISBN: 0-612-88046-X
Our file Notre référence
ISBN: 0-612-88046-X

University of Alberta

Library Release Form

Name of Author: Ehud Sharlin

Title of Thesis: Blurring the Boundaries: Spatial Tangible User Interfaces for Cognitive Assessment

Degree: Doctor of Philosophy

Year this Degree Granted: 2003

Ehud Sharlin
48 Gilboa St., Alfe-Menashe
Israel, 44851

Date: July 21, 2003

University of Alberta

Faculty of Graduate Studies and Research

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled **Blurring the Boundaries: Spatial Tangible User Interfaces for Cognitive Assessment** submitted by Ehud Sharlin in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

Ben Watson
Dr. Benjamin A. Watson

Jonathan Schaeffer
Dr. Jonathan Schaeffer

Dr. Herbert Yang

PR
Dr. Lili Liu

PR
Dr. Albert A. Rizzo

Date: 14 JULY 2003

אמר רבי עקיבא:

" 'ואהבת לרעך כמוך', זה כלל גדול בתורה"

Said Rabbi Akiva:
" 'Love Thy Neighbour as Thyself', this is a fundamental principal of the Torah"

"כל העולם כולו גשר צר מאוד. והעיקר, לא לפחד כלל"

רבי נחמן מברסלב

"The whole wide world is but a very narrow bridge. The essence is not to be afraid at all"
Rabbi Nachman of Bratslav

Abstract

Human computer interaction (HCI) research is a multi-disciplinary effort struggling to study and remove the barriers between people and the computers they use. As computer technology augments more and more of human activity, HCI research becomes ever more essential. A crucial aspect of HCI research is directed at the exploitation of our innate tactile and spatial abilities. An extremely successful outcome of this effort is the common mouse. Recently, the notion of a tangible user interface (TUI) has emerged, suggesting more elaborate use of physical objects as computer interfaces. Our work represents one of the very first attempts to move TUIs beyond conceptual prototypes into meaningful applications.

The uniqueness of TUIs lies in their spatiality. We discuss a subset of TUIs which we define as *spatial TUIs*. We propose and use a practical set of heuristics for developing good TUI applications and for evaluating existing ones. Following these heuristics, we designed and tested two novel TUIs for cognitive assessment: *Cognitive Cubes*, for assessing three-dimensional (3D) constructional ability, and the *Cognitive Map Probe* (CMP) for assessing cognitive mapping ability. In testing, both TUIs were confirmed to be sensitive to factors known to affect cognitive performance. Cognitive Cubes and the CMP are the first systems to automate fully these 3D neuropsychological tests, increasing the potential for better assessment's consistency, flexibility, reliability, sensitivity and control. Cognitive Cubes and the CMP are the first experimentally tested 3D TUI applications.

Acknowledgments

For the last five years I have been supervised by, and have enjoyed the collaboration and friendship of some really wonderful individuals. This work became possible thanks to them.

I was extremely privileged to be one of Ben Watson's students. In Ben I have found the qualities of a true mentor and teacher, a wise, visionary and open-minded researcher, and a down-to-earth and hard working professional. Ben's brightness and enthusiasm initiated my work and constantly pushed me forward. I am very grateful for the intensive and fruitful relationship we shared, and for the huge impact Ben had on my work and on my life. I was also extremely fortunate to have Jonathan Schaeffer as a mentor. Jonathan practically brought me to the University of Alberta, and was always there for me with an advice and a helping hand. Co-supervising my work, Jonathan's wisdom, clear judgment and guidance, and vigorously critical eye had great influence on this research as well as on my life.

Lili Liu supported my work from the very first stages, as a member of my supervisory committee, with her penetrating insight, professionalism, ability to bridge disciplines, to collaborate and her endless positive energy. My good friend and colleague Steve Sutphen, a true modern Renaissance man, supported me in almost every aspect of this work, and in many aspects of my life. I am going to sorely miss the long nights of thinking, working and cross-country skiing we shared.

I am indebted to my external examiner, Albert "Skip" Rizzo, whose vivid research vision guided me throughout this work. Herb Yang helped me generously with hardware and with detailed and useful feedback as a

member of my supervisory committee. My good friend and colleague Pablo Figueroa was an important part of my life and work, I am going to miss you Amigo. Good luck to you wherever you are heading.
My work was possible thanks to the great help and collaboration of a truly amazing group of researchers: Yoshifumi Kitamura, Yuichi Itoh, John Frazer, Robert Lederer, Adrien Cho, and members of their research groups. Thank you all so much.
I am very grateful for the collaboration, help and friendship of Akihiro Kishimoto (Kishi), Amr Kamel, Andreas Junghanns, Andrew Nisbet, Andrzej Zadorozny, Anup Basu, Ayman Ammoura, Barb Johnston, Ben Ragosin, Benny & Barby Mer, Dana Cobzas, Darren Rorabeck, Darse Billings, Dima Brodsky, Doug Weir, Ed Cornell, Edith Drummond, Hong Zhang, Jonathan Molcho, Julian Fogel, Leeor Kronik, Lorne Webber, Mario Costa Sousa, Mark Green, Martin Jagersand, Minglun Gong, Morten Fjeld, Natan Merom, Oscar Meruvia, Osmar Zaiane, Peter Woytiuk, Roel van der Goot, Roman Fedoriw, Ronald Pang, Russ Greiner, Ruyam Acar, Shaojun Zhao, Stan Cabay, Tamer Ozsu, Tim Van Allen, Uzi Koren, Xuejie Qin, Yanxia Jia, Yehuda Nachshon, Yngvi Bjornsson, Yossi Mashiah and many others.
I am grateful for the endless support and love of my mother, Meira-Mer Sharlin, my father, Avraham Sharlin, my brother Elad, my sisters Efrat and Adva and my grandfather, Nahum Mer, back home.
Thank you and sorry for all the time I spent working instead of playing with you, my beloved and wonderful Noga and Naomi. Thank you Amira, my love.

In memory of Leb (Arie) Mer and Hiroki Tokashiki (Toka).

Table of Contents

List of Figures

List of Tables

Chapter I
Introduction

1.1 Designing Effective Tangible User Interfaces

The massive acceptance of technology in the developed world has created a considerable shift in attention from the algorithms that make computers tick, to the interfaces that enable computers and humans to communicate. This attention created the relatively new domain of human computer interaction (HCI). HCI can be simply defined as "the study of people, computer technology and the ways these influence each other" [33].

The last 20-30 years have seen rapid technological progress that has dramatically affected both hard-core computing technologies (such as processing power and storage size) and various user-centered technologies (such as graphics and display capabilities, sensing and tracking). At the same time, wide consumer assimilation kept these technologies affordable. From these processes emerged various pioneering HCI disciplines, such as speech and handwriting based interfaces, and virtual reality (VR). Tangible user interfaces (TUIs) are a newly emerging HCI discipline.

Like all HCI research domains, TUIs are constrained and shaped by the fact that computers are digital entities while their users are human beings. TUIs attempt to bridge this gap by exploiting the innate relationships between people and physical objects. The common mouse is an example of the huge impact this approach can have on HCI (we discuss our thoughts of the mouse as a TUI in Section 2.4).

Fitzmaurice et al. were the first to distinguish TUIs, or graspable user interfaces as they termed them initially, from other interfaces [39]. Fitzmaurice defined a graspable user interface as: "a physical handle to a virtual function where the physical handle serves as a dedicated functional manipulator" [40]. Ishii and Ullmer, who suggested and established the term TUIs, define them as "devices that give physical form to digital information, employing physical artifacts as representations and controls of the computational data" [162]. Both of these definitions highlight the coupling between the physical object and the virtual information or function it embodies as the essence of a TUI.

Our work builds on these general TUI definitions, enhancing them with another distinctive layer. We believe that the coupling of physical and virtual entities is not sufficient for a TUI application to be useful and meaningful. Simply put, we believe that this coupling must make sense spatially; or more rigorously, follow what we call an intuitive spatial mapping. We define spatial mapping as *the relationship between the object's spatial characteristics and the way it is being used.* We focus our attention on spatial TUIs, a subset of TUIs that embody higher spatial expressiveness, and thus the ability to allow more intuitive spatial mapping between the physical object and the virtual task. We define spatial TUIs as *tangible user interfaces used to mediate interaction with shape, space and structure in the virtual domain.*

This dissertation formalizes our thoughts into heuristics of effective TUI design. This set of rules is then demonstrated in practice, by designing and testing meaningful TUI-based tools. The reader should not overlook an inherent limitation of our approach. We do not attempt to demonstrate that TUIs are useful if and only if our heuristics are followed. However, we do point to several instances where less spatial TUIs failed to show effectiveness as possible evidence.

1.2 Thesis Statement and Contributions

Our thesis is that:

> *Tangible user interfaces can be feasibly used in meaningful computer applications.*

In order to confirm this thesis the dissertation includes the following contributions:

1. *A set of heuristics for identifying and developing good TUI-based applications, and for evaluating existing ones.*
2. *A taxonomy for TUIs and a thorough review of the state of the art, based on our heuristics.*
3. *Design and implementation of two TUIs aimed at cognitive assessment tasks, acting as a successful manifestation of our heuristics: Cognitive Cubes for assessing 3D constructional ability and the Cognitive Map Probe (CMP) for assessing cognitive mapping ability.*
4. *An experimental testing of our TUIs. This testing is among the very first applied to meaningful applications of TUIs, and the first experimental testing of 3D spatial TUIs.*[1]
5. *Spin-offs of our efforts are the first automatic system for neuropsychological assessment of 3D constructional ability*

---

[1] Rasa (Sections 2.5.2 & 2.12.4) and Senseboard (Sections 2.8.4 & 2.12.5) are two-dimensional (2D) TUI applications that were successfully validated in testing at roughly the same time.

(Cognitive Cubes), and the first detailed automatic system for neuropsychological assessment of cognitive mapping ability (the CMP).[2]

1.3 Our Quest

We started our work in 1998-1999 hoping to design and build the very first 3D spatial TUI (these plans are summarized in [138]). During this period we began developing our own thoughts about what makes a useful TUI, emphasizing the need for intuitive spatial mapping between the TUI's physical properties and its task (see Sections 2.2-2.3). We later became aware that development of similar 3D TUIs was being pursued by at least two other major industrial groups, which already had working prototypes and were about to publish. We also discovered the pioneering work of Frazer [51-55] and Aish [1,2] in the early 1980s, work that took much of the research novelty out of these more recent technological achievements (see Section 2.11 for a detail discussion of these projects).

Given these developments, we believed that designing the next generation of 3D TUIs would be a major electro-mechanical engineering achievement, but would hardly constitute a Ph.D. in Computing Science. On the other hand, we had already gained a good knowledge of the field and developed our own philosophy on how TUIs should be used. We were aware that current TUIs were usually portrayed as "cool" ideas, with very little practical value and with a marked lack of user-based testing and

---

[2] We are aware of a few previous instances in which automatic techniques were used to support assessment of cognitive mapping or wayfinding ability (Chapter 3). However, the CMP is the first automatic system to offer a detailed automatic assessment, with each action recorded and analyzed, not just the end result.

confirmation. We therefore decided in 1999-2000 to shift our attention to finding effective TUI applications.

Our search for potential TUI applications was based on the following thoughts:

1. Design of TUI applications should follow our suggested set of heuristics (Section 2.2) as closely as possible. Our first and foremost rule-of-thumb: there should be an intuitive spatial mapping (see definition in Section 2.2.2) between TUI and task.
2. TUIs need to develop beyond conceptual prototypes into meaningful and practical applications.
3. Given the current infant state of spatial TUI development, the chosen applications should require only a simple level of expressiveness, without compromising the significance of the resulting system.

Our search for meaningful applications took several directions, most profoundly an ethnographic study of neuropsychological assessment and training tools (see Chapter 3) and of acute orientation and mobility needs of the visually impaired (see Chapter 6). Both domains offered a need for spatial, tangible interaction with computers and both domains did not require a high level of spatial expressiveness during such interaction (that is, the required level of detail, miniaturization and spatial flexibility was not high). On the contrary, the low level of detail they required was actually seen as a benefit.

We believe that our efforts successfully demonstrated the feasibility of TUIs in meaningful computer applications. As the title of this dissertation hints, we see our TUIs as a step towards further blurring of the existing boundaries between the digital domain and the physical one, acting more as "a tool" rather than an interface to a computer.

1.4 Organization

Chapter 2, "Tangible User Interfaces", is an extensive literature review of this emerging field. We believe this to be the most up-to-date and complete review available. We begin by exploring the relationship between humans and physical objects, and from that extract a set of heuristics for identifying and developing good applications of TUIs. These heuristics are used throughout the chapter as evaluation criteria of the various TUI systems we reviewed. We also present our own approach for a TUI taxonomy, spatial TUIs, and a discussion of previous attempts to test TUIs and empirically compare them to other interfaces.

Chapter 3, "Cognitive Assessment", is a brief introduction to the cognitive assessment domain. The chapter concisely covers various neuropsychological issues that are later used in our work, such as constructional ability, cognitive mapping ability, mental rotation test, assessment automation, the elderly and Alzheimer's disease (AD). These important matters are merely presented in summary to familiarize the reader with the concepts and terminology used in later chapters.

Chapter 4, "Cognitive Cubes", and Chapter 5, "The Cognitive Map Probe", are detailed presentations of our TUI-based 3D constructional ability assessment system (Chapter 4), and cognitive mapping ability assessment system (Chapter 5) and the experiments that we conducted to test them.

Chapter 6, "Conclusion and Future Work", presents a summary of our contributions, a detailed discussion of proposed future work, including a test comparison for the CMP, a TUI-based orientation and mobility application for the visually impaired, and some concluding thoughts.

Appendices A and B detail experimental issues, such as consent forms, information letters and experimental protocols, for Cognitive Cubes, and the CMP, respectively. Appendix C lists the abbreviations used in this dissertation.

1.5 Publications

Chapter 4 is closely based on our Cognitive Cubes ACM-CHI publication [137]. The early part of Chapter 5 is loosely based on several CMP-related publications [21,140-142,154].

Chapter II

Tangible User Interfaces

2.1 Scenario and Motivation

> Picture yourself designing a prototype, say an extension to your veranda or a small yacht, using a 3D CAD (Computer Aided Design) tool, but instead of mouse, keyboard and screen you are using a set of small scale objects, that match in shape, color and texture your design needs. As you assemble and connect the small objects a corresponding digital replica is being created on your PC, following in real time the progress of the physical design. After you are content with the physical model you assembled you use a set of physical tools to evaluate your design. Using a drill-like device your run a complete strength analysis on your model. The model then reflects the digital outcomes by physically coloring weak parts that need to be altered. Using a physical eraser-tool you confirm the removal of some of the frailer subparts. As you perform these changes the physical model transforms, disconnecting some of the objects and squeezing or expanding others, simultaneously updating the strength analysis color output, closely following your edits. After making several changes you examine the physical model again and finally decide to print several 3D copies of it.

Science fiction? Probably, but maybe not for long. In this chapter we present a technical survey of the state of the art of TUI technology. We cover the pioneering work performed in this field more than 20 years ago on one hand and the exciting, recent developments in the field on the other. For our presentation we employ human-perception and human-activity related criteria for TUI evaluation. Although we attempt to overview the entire TUI domain, our presentation is guided by our belief that TUIs are most valuable in spatial applications. We try to emphasize what we believe are important attempts to move TUI research beyond conceptual prototypes into applications.

We begin this chapter by introducing TUIs as an HCI paradigm that evolves from the relationship between humans and physical objects. From this relationship we extract a simple set of observations that will be used throughout the chapter as heuristics for evaluating the merit of different TUIs (Section 2.2). We then examine TUI definitions, including our own thoughts on spatial TUIs (Section 2.3) and take a close look at the common mouse as a TUI (Section 2.4). Following that, we cover information containers (Section 2.5), the use of real objects as TUIs (Section 2.6) and Waldo interfaces (Section 2.7). We discuss in detail the tabletop TUI paradigm (Section 2.8), review instances of spatial TUIs explicitly designed for input of vectors, curves, surfaces and volumes (Section 2.9), spatial TUIs for 2D topology input (Section 2.10), and spatial TUIs for 3D structural input (Section 2.11). We summarize the chapter by detailing efforts to experimentally test and empirically compare TUIs to non-TUIs (Section 2.12), briefly covering closely related work outside the TUI domain (Section 2.13) and by presenting some concluding thoughts, including a discussion of possible future trends in TUIs and a review table (Table 2.1, on page 77) detailing the TUI instances presented in the chapter (Section 2.14).

2.2 Introduction

2.2.1 Why TUIs?

Our natural interaction with the world relies, among other abilities, on our innate ability to manipulate tangible objects. From a very early age we can naturally move, manipulate, assemble and disassemble a seemingly endless variety of physical objects with very little cognitive effort (for example, a child is able to build physical sandcastles, snowmen or Lego buildings, constructing elaborate structures in a physical 3D space). Yet, performing similar tasks in a 3D computer-based virtual world using the WIMP (windows-icon-menu-pointer) interface (from here on we will refer to it as the standard user interface or the standard UI) can become extremely cumbersome and complex. We, like others, believe that a large part of this difficulty results from the failure of current computer interfaces to fully engage our innate spatial abilities.

What formalized explanation can we use to clarify the ease of building models or manipulating objects in the real world, as opposed to the hardship involved in performing similar tasks using the standard UI? Based on previous research, we consider three explanations: the support of intuitive spatial mappings, the support of unification of input and output (or I/O unification) and the support of trial-and-error actions. As we will detail later in this section, we adopt these three explanations and use them as near-orthogonal criteria in our evaluation of different TUIs. We believe that these proposed criteria are simple heuristics for TUI examination and evaluation. Let us take a closer look at these explanations.

2.2.2 Spatial Mapping

When we interact with objects in the real world we usually do not have to consciously apply complex thought in order to manipulate or use them. Their function and use is inferred from their qualities: shape, weight, size, etc. This functionality, expressed through the object's physical form, is called "affordances" [43,56,106]. Following, we can say that many physical objects afford their behavior, that is, their physical form reveals a clear representation of their functionality. In a similar fashion, Norman discusses a "natural mapping" between an object and its functionality. In the case of clear natural mapping, an object's functionality is unmistakably expressed in its physical shape by "taking advantage of physical analogies and cultural standards" [106]. Natural mapping can be viewed as "primitive" since it materializes our innate schooling as spatial beings. Once we are familiar with the spatial vocabulary of an object, we can use this knowledge to translate our spatial skills to new tasks, essentially forming new spatial mappings for the object (for example, a child's initial use of a spoon for picking up food and only later for mixing fluids).

When thinking of spatial mapping we are also informed by Beaudouin-Lafon's "degree of integration" and "degree of compatibility" properties [12]. "Degree of Integration" considers the *ratio between the number of spatial degrees of freedom (DOF) of the object's function and the object.* When we interact directly with physical objects in the real world their degree of integration naturally equals to one since there is no separation between the object and its function. Much like the notions of affordances and natural mapping, "degree of compatibility" *measures the similarity between the actions performed on the object, and the object's response* [12]. This concept can be clarified by examining a car's steering wheel. If turning the steering wheel to the left turns the car to the left then there is high similarity between action and response, and high degree of compatibility. However, if turning the steering wheel to the left turns the car to the right there is low similarity between action and response, and low degree of compatibility.

We use all these concepts when discussing spatial mapping, which we define as the relationship between the object's spatial characteristics and the way it is being used. Note that unlike taxonomies that refer only to input or only to output (see, for example, [41]), spatial mapping is concerned with the relationships between the two.

The physical world usually offers clear spatial mapping between objects and their functions. In the digital world most user interaction techniques, particularly in 3D modeling, include a set of rules and controls that manipulate their functionality. However, these rules and controls, implemented with the restrictions of the standard UI, are far from enabling an intuitive spatial mapping between interface and application.

Being physical, all TUIs are naturally spatial (an even stronger claim was made by Fitzmaurice et al. who consider all input devices as spatial samplers with varying abstraction [41]). However, the ever-spatial TUI can be matched to various applications of varied spatiality. We argue that the quality of the resulting spatial mapping can vary dramatically and plays a crucial role in the success of the applied TUI. In the more obvious cases, a TUI can be used for an application that is not intrinsically spatial (such as a database). In such cases we argue that the spatial mapping between interface and application will be unintuitive, or difficult to perceive. We argue further that in such cases it might be hard to find a real benefit from using a TUI over the standard UI.

Even in cases of applications which are fundamentally spatial, the number of spatial DOF the TUI and task offer can vary, potentially causing the degrees of integration to divert from the optimal value of one. When a TUI and a task match in the number of spatial DOF, the TUI's spatial mapping quality still depends on the affordances embodied in the TUI and how well it maps to the task. For example, an alarm clock that with a sizeable, protruding, on-off alarm switch affords its task (quickly turning off a noisy alarm in poor lighting conditions) much better than an alarm clock with a small and hidden switch, although both switches have exactly the same single spatial DOF.

We believe that the art of reaching an intuitive, finely tuned, spatial mapping between a TUI and an application is probably the most important part of TUI design. To this end, the first rule of thumb that TUI design should follow, in order to reach intuitive spatial mapping, is to maintain a good degree of integration (that is, a value close to one), matching the number of DOF between the TUI and the task. The second rule of thumb is to preserve a high degree of compatibility, keeping high similarity between the action performed using the TUI and its results. Special care should be taken with the mechanism that maps the TUI's DOF to the task's DOF (assuming that the degree of integration is equal to one). Ideally, the dynamics underlining the different DOF should match. For example, rotation of the TUI should be mapped to a rotation movement in the task domain, linear TUI movement to a linear movement in the task domain, etc.

Last, comes the crucial, and less quantifiable, notion of affordances and natural mapping. How does one take advantage of interactions between physical objects and applications that are "hardwired" in users' behavior and cognition? How do we keep the spatial mapping between TUI and task instinctual and familiar? Answers will vary according to application specifications and users ethnography. To emphasize the complexity of this design parameter let us briefly look at the intuitive use of physical objects at different ages: in infancy physical objects can innately be picked up, handled and assembled; in elementary school physical objects are used as a natural means for writing and in high school as an intuitive means for controlling a car. Although we naturally use physical objects as interfaces for physical tasks, we are also trained in using them for more abstract tasks (for example, flipping a coin to make a decision, or calculating with an abacus).

We argue that spatial mapping will always benefit by exploiting skills gained over a lengthy training period. Physical objects that are used from an early age for a certain spatial task would eventually appear natural for that task. We believe that TUIs that take advantage of such a primitive spatial mapping will be more effective than others that do not. While we find it hard to clearly state which tasks will or will not benefit from the use of physical objects, we believe we can state which interfaces might be beneficial in certain tasks, following the intuitive spatial mapping rules stated in this section.

2.2.3 I/O Unification

Interaction with physical objects naturally unifies input and output. We see two components to this natural unification: the coupling of action and perception space and the clarity of state.

When we interact with an object in the real world, our hands and fingers (parts of our action space) coincide, in time and space, with the position of the object we are handling (part of our perception space) [43,46]. This spatial and temporal coupling of perception space and action space focuses attention at one time and place, and enables us to perform complex tasks. Yet the standard UI separates mouse from screen, input from display, and action from perception, dividing attention and requiring mental mapping of one space to another. Since TUIs are physical objects, they offer tangible, haptic feedback, unifying perception and action in the tactile domain. In the visual domain, input and display surfaces are often identical, strengthening this unification. Moreover, the objects in a TUI typically represent only one virtual object (unlike the mouse), bringing perception and action into still closer agreement.

The state of physical interfaces in the real world is usually clearly represented. For example when we tie or untie our shoelaces their tactile stretch, physical pressure on the foot and visual appearance will clearly represent their state and inform us if we are progressing toward our goal. In HCI this is not a given and the state of the application and the UI do not necessarily mirror each other. In the standard UI realm the need to clearly express the system status is well established (see for example Nielson's visibility of system status [105] or Shneiderman's classic direct manipulation [145]). However, the standard UI extends this notion only to the display which carries the sole responsibility of reflecting the application state as soon as it is altered by the input device. As the identification between physical input elements and application elements grows so does the expectation that the state of the input devices should also mirror the application state. In TUIs, this expectation is extended not only to the visual state but also to the physical state of the input device.

2.2.4 Support of "Trial-and-error" Actions

When we build a physical 3D model, we actually perform an activity that is both cognitive, or goal related, and motorized [43,46]. Such a physical task involves both pragmatic and epistemic actions [40,43,46,77]. Pragmatic actions can be defined as the straightforward maneuvers we perform in order to bring the 3D shape closer to our cognitive goal. Epistemic actions, on the other hand, use the physical setting in order to improve our cognitive understanding of the problem. Some of these epistemic maneuvers will fail and will not bring us any closer to our goal, while others will reveal new information and directions leading to it. In fact, this information might have been very hard to find without trial-and-error [77]. Please note that for clarity we choose to discuss the support of trial-and-error actions, rather than use the less familiar notion of epistemic actions.

Kirsh and Maglio used the game of Tetris as an empirical example [77]. Tetris is a straightforward and simple goal-oriented computer game in which the player is trying to fit different geometrical shapes falling from the top of the screen into horizontal rows. Common action in Tetris is to rotate the falling shape around its axis before attempting to fit it into the horizontal row. Kirsh and Maglio claim that this rotation is an easy exploitation of the Tetris world in order to reveal new cognitive information ("How will this shape fit?"), rather than performing a complex cognitive action (mentally rotating the shape). Although physically rotating the shape doesn't serve an immediate pragmatic goal, it reveals important new information through simple trial-and-error in the world.

Physical 3D modeling, in most of the simpler cases, provides us with both pragmatic and epistemic tools for performing our tasks, with a low cost for mistakes made while executing an epistemic step [43]. While building models with Lego blocks we almost inattentively test the validity of certain actions, some of which do not lead directly to our goal. The price of correcting an action that proves to be erroneous is minimal. On the other hand, the standard UI is geared towards pragmatic actions. The idea of simple, primitive epistemic tools does exist, but for complex tasks such as 3D modeling the usability of these epistemic tools is substantially inferior to their physical world parallels. For example, the "undo" operation is linear, meaning that to "undo" a single erroneous operation, you have to also "undo" all the operations that followed it.

Furthermore, we believe that trial-and-error activity is better supported when the user can handle the interface internal state via multiple points of access. Many TUIs support this premise, having multiple physical objects acting as persistently coupled mediators to the application objects. Without persistent coupling the user is left with a history of coupling and decoupling between a few physical objects and many application objects (for example, in the standard UI a single physical object, the mouse, is coupled and decoupled to any number of application objects). The absence of persistent coupling naturally leads to a linear “undo” metaphor that poorly supports random, trial-and-error operations.

2.2.5 Conclusion

To conclude, TUIs make sense since they engage our natural talents for handling every-day objects in the physical world. As we will show later, although different TUIs are variously successful in achieving this engagement, they all perform better than the standard UI in at least some of above criteria.

2.3 What Makes It a TUI?

The use of physical objects as means of interaction with computers is far from being new. Early examples date back to 1955, when Bert Sutherland developed the SAGE system, using the first light pen for interaction [19]. Later, in 1963, Ivan Sutherland designed Sketchpad [149,172] which enabled the user to input simple curves on a screen using a light pen. Doug Engelbart invented the first mouse (Figure 2.1) in 1963 [37]. Today’s common user interfaces originate from this basic research.

Dourish [35] sees TUIs as part of a larger HCI movement (which he

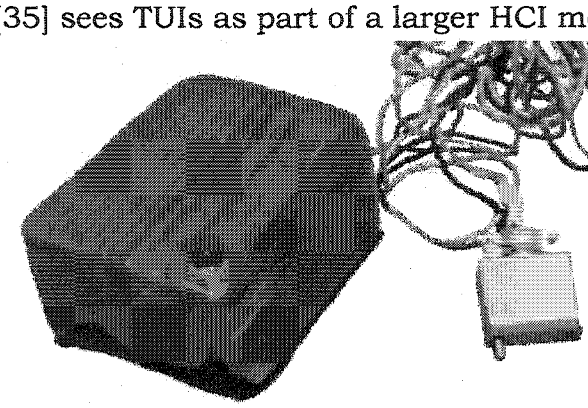

Figure 2.1 - The first mouse

calls "Embodied Interaction") that attempts to capture the way people experience the real world. Hence TUIs shouldn't be viewed only as some new physical interaction metaphors, but as interfaces that draw strongly on the way they are being used in the everyday, non-computerized world. Ullmer and Ishii, from the Tangible Media Group at the MIT Media Lab, define TUIs as "devices that give physical form to digital information, employing physical artifacts as representations and controls of the computational data" [162]. Ullmer later defines TUIs as "systems that use spatially reconfigurable physical artifacts as representations and controls for digital information" [160]. Ullmer and Ishii state the following four TUI key characteristics [162]:

1. *Computational coupling between the physical representation and the digital information.*
2. *Perceptual coupling between the physical representation and the digital information.*
3. *The physical representation affords interactive control.*
4. *The physical representation embodies, to some extent, the digital state of the system.*

While not challenging these definitions of TUIs, we believe that in order for TUIs to shine through as significantly more useful than the existing standard UI parallels, they have to offer more than strong coupling between physical and digital representations or, more explicitly, they have to offer intuitive spatial mapping between TUI and application (see Section 2.2). We argue that intuitive spatial mapping is more evident in a subset of TUIs that manifest higher level of spatial expressiveness. We define this subset as *spatial TUIs,* which are *tangible user interfaces* *used to mediate interaction with shape, space and structure in the virtual domain.*

2.4 The Mouse as a Test Case

Is the mouse a TUI? The mouse was the first common UI that captured some of the human spatial senses and abilities while interacting with computers. We can attribute its success, and more importantly in this context, support for the standard UI, to its engagement of the human spatial-tactile abilities. Given the mouse's substantial success, it is interesting to examine the mouse's qualities in regard to the TUIs definitions of Section 2.3 and our heuristics for tangible interaction of Section 2.2.

The mouse is probably not a TUI according to Ishii and Ullmer's general definition (Section 2.3). While it is a physical object used to control computational data, it is hardly a means of giving physical form or representation to digital information. The mouse's power is in its generality. In most of its tasks, most of the time it is not mapped to a unique digital entity but rather is attached and detached constantly to different entities, according to need. It hardly ever represents these digital entities, but rather is used to control them for a certain period of time. Let us look at the mouse by our Section 2.2 heuristics. The mouse as a 2D physical entity is highly spatial. It is fairly easy to hold the mouse and roll it on a surface. The mouse has a generic physical form that allows the user to grasp it with relative comfort [40]. We can claim that for limited durations the mouse is spatially mapped extremely well to certain 2D tasks, from handling the standard UI (for example, moving a cursor, selecting and moving windows), to the numerous planar-data editors (for example, word processors, spreadsheets, etc.).

At the same time the mouse is used for numerous other, not essentially 2D, interaction tasks that range from changing the viewpoint of a virtual character in a shoot-em-up PC game to editing a detailed 3D structure using CAD software or editing multimedia sequences. Many of these tasks are highly spatial, while others are less spatial or involve more spatial dimensions than the 2D that the mouse affords. In many of these tasks the mouse's clear and natural spatial mapping, that was so evident in the simple 2D tasks, breaks down. The user is left with a generic but rather non-intuitive interface, and with the need to compensate for the inadequate spatial mapping by application-specific training. To make things worse, in many instances the mouse continuously refers to different entities and is attached and detached to and from them as the task progresses. The mouse does not support I/O unification: action space (the hand moving the mouse) and perception space (the cursor moving on the screen) are separated and the state of the input device does not reflect the application state. As a pointing device the mouse directly supports goal-oriented actions, but in a vague way the mouse also supports some trial-and-error actions. For example, in the WIMP environment the user can arrange several windows across the desktop and then change their location in a random manner with no relation to the linear order in which they were placed originally (a non-linear "undo"). However, in most standard UI applications it is arguable what information, if any, will be revealed by these trial-and-error actions.

We can conclude by highlighting the fact that in spite of offering much less than ideal TUI qualities, the mouse was the first common UI that actually made direct use of the user's spatial-tactile abilities. For many of its everyday WIMP applications, the mouse does offer an intuitive and strong spatial mapping. The mouse's great success and the fact that it is currently part of the standard UI suggest the huge potential of more ideal TUIs.

2.5 Information Containers

2.5.1 Simple Instances

There are a number of applications that use physical objects as information containers. A classic example is Bishop's Marble Answering Machine (Figure 2.2), which used marbles as tangible representations of incoming voice messages. The user could pick up a marble and place it in a certain notch on the machine in order to listen to the message embodied in it, or put it in a different place in order to dial the caller's number [26,69].

Simple information containers were recently commercialized with two musical toys: Neurosmith's Music Blocks [103] and Lego's Music Builder Composer [88] (Figure 2.3). The two quite similar interfaces enable a child to compose music by placing tangible objects on a board, literally "building music". Objects can represent musical instruments,

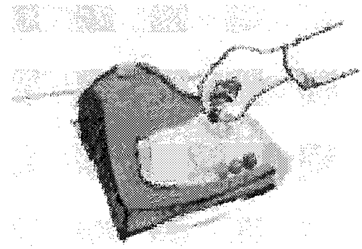

Figure 2.2 – Bishop's Marble Answering Machine

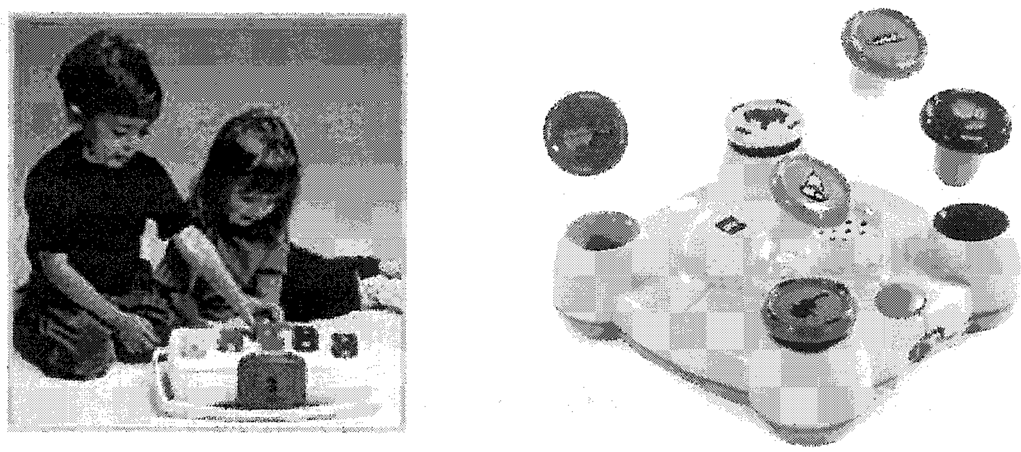

Figure 2.3 – Music Blocks (left) and Music Builder Composer (right)

rhythms, animal voices, etc. The interface synthesizes the instruments the child places into a musical piece that changes interactively as objects are inserted and removed from the board.

In mediaBlocks (Figure 2.4), Ullmer et al. used wooden blocks with an embedded electronic ID as physical icons [163]. The physical icons appeared as if they contain data, and supported data transport and simple data sequencing by ordering the blocks on a rack. In a similar manner, Logjam was designed for a media editing—“video logging”—application, or the process of “finding and marking locations, events and behaviors in a video sequence” [23]. Logjam was based on a Scrabble-like 4×12 rack and a set of wooden blocks, each containing, or representing, a different video event category. The users could place the category blocks on the rack according to the events being watched on video in order to correctly categorize them. While the user is expected to treat the tangible information containers as if they actually embody data, the containers do not really hold any data other than their ID, which is pointing to the “real” data item.

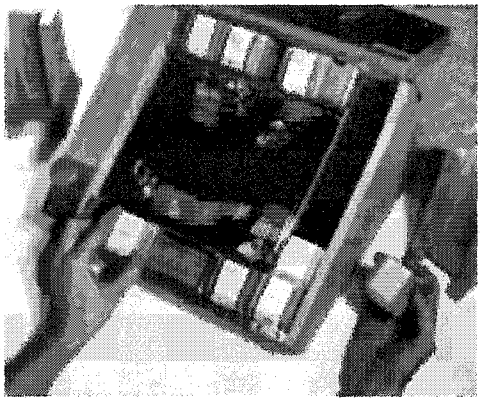

Figure 2.4 - Ullmer's mediaBlocks

Let us try to look at simple information containers using our TUI heuristics. Information containers can potentially offer intuitive spatial mapping by matching their spatial shape to the type of information they embody. Another way to achieve natural spatial mapping is by matching the spatial location in which the information container becomes active to its application. For example, Ullmer's mediaBlocks can "paste" a printout when placed near a printer, or a video sequence when placed near a projector. As was demonstrated in mediaBlocks and Logjam, a simple physical constraint, like a rack, can be used to spatially map the information container's location to a simple set of logical rules, such as ordering (for example, being closer to the right of the rack means you are higher in hierarchy).

These information containers do not support I/O unification. Actions performed with an information container generate an effect that is external to the TUI (for example, voice through a speaker in the marble answering machine, images on a screen in mediaBlocks). Although the user can perform simple and direct actions using information containers, it is hard to imagine how they support trial-and-error activity.

Information containers' strongest quality is their ability to embody digital data in a physical form, but we argue that this does not necessarily make them a strong TUI instance. CDs or diskettes also contain information, but do they offer natural, instinctual spatial mapping? Are they TUIs? We do not think so. An information container can build on spatial, visual and tactile similarity between itself and the information it contains to exploit users innate abilities and offer some natural spatial mapping. However, in cases where the information container is using a novel mapping between its spatial characteristics and the information it contains, we think that the natural spatial mapping advantages it offers are arguable at best.

2.5.2 Paper Based TUIs

A number of interfaces to the digital domain were designed around simple paper notes, probably the most fundamental examples of pre-digital information containers. McGee et al. underline some of the powerful qualities of paper as an interface [96]. Paper is lightweight, extremely cheap, supports high resolution information (handwritten or printed) and is quite reliable (paper will not crash or lose information when the power is down).

Wellner's pioneering work [104,173,174] used physical paper documents as interfaces. His DigitalDesk Calculator augmented paper with electronic data employing an over-the-desk camera and projector.

McGee's Rasa [96] is a paper-based interface for a military command post application. The interface augments an existing command post tactical infrastructure, a constantly updated paper map with handwritten sticky-notes, each representing a military unit and its related information. Instead of forcing the users into a new interface, Rasa uses the reliable and familiar paper-based infrastructure by sampling both the physical content of the sticky-notes (using a tablet to digitize the handwriting), their spatial location on the map, and the verbal comments or commands associated with them (using voice recognition). Through the sticky-note interface Rasa can maintain a full, up-to-date, digital version of the tactical map with all its information layers. We will mention Rasa again in Section 2.12.4 when we examine a user study performed with the system, comparing the augmented paper TUI to plain paper sticky-notes.

Mackay et al. present an impressive ethnographic work on the use of paper flight strips by air traffic controllers [95]. A flight strip (Figure 2.5) is a band of paper that is used extensively by air traffic controllers around the globe, each representing a flight and its specific data and condition. Flight strips can embody information in a number of ways: they can be printed or written on, placed on a rack or held up silently in the air for immediate attention of the other members of the controllers group. Mackay et al. claim that air traffic controllers remain loyal to the

Figure 2.5 – Air traffic controllers interacting with paper flight strips

paper flight strip and reject many suggested WIMP replicas, not due to fear of technology, but rather because of the tangible, spatial and visual superiority of paper. Much like McGee et al. they argue, while making the first steps towards a prototype implementation, that the right way to introduce automation to this domain is by "automating" the paper flight strips and using them as TUIs: capturing their content and augmenting them (using computer vision and projection, respectively).

Paper is well established as an interface of choice in many human activities. We believe that automation of such activities by exploiting paper as a TUI infrastructure can maintain the strong intuitive spatial mapping between the paper and the task. In some of these applications using paper as a TUI can potentially support trial-and-error activity and I/O unification. For example, an interface like Rasa could enable users to spatially place paper-interfaces on a map in either a direct, goal driven way, or in a trial-and-error manner. Furthermore, if the interface augments the paper with digital information (for example, using projection) the TUI can also afford high levels of I/O unification.

2.6 Real Objects as TUIs

Straightforward use of real-world objects and props as TUIs can sometimes simplify very complex interaction tasks. A classic example is Passive Real-World Interface Props that were developed along these lines for a neurosurgical visualization application in 1994 [40,67]. Physicians were facing difficulties while attempting to manipulate and define slices of volumetric head scans. The Passive Real-World Interface suggested a simple method for this interaction. A doll's head and a plastic plate were both tracked in 6 DOF while the scan data and slice plane were being

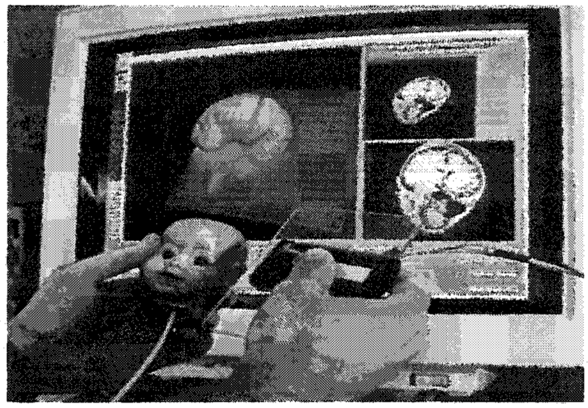

Figure 2.6 – Hinckley's Passive Real-World Interface

mapped to them. The user interacted with the volumetric samples by simply positioning the plate beside the doll's head defining the desired slice orientation (Figure 2.6). Each positioning resulted in a simple definition of the slice plane. This intuitive spatial mapping between the physical doll and plate, and the volumetric slice, drastically simplified a task that can be quite complex when performed with the standard UI.

Along the same lines of simplicity through intuitive spatial mapping, Tonka® Workshop (Figure 2.7) is a simple and low-cost toy [84,159]. Nevertheless, it is a perfectly good example of a mature, task oriented TUI. The interface is a set of tangible workshop hardware

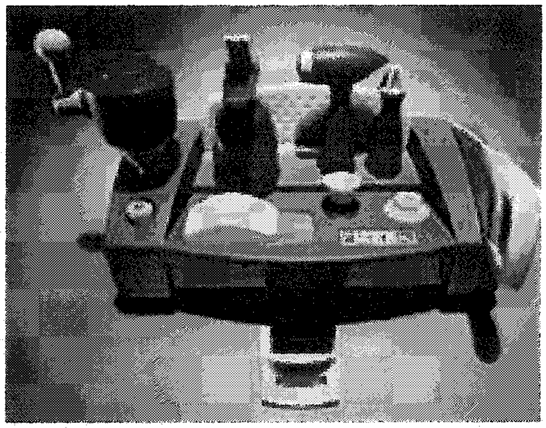

Figure 2.7 – The Tonka® Workshop interface

consisting of a hammer, screwdriver, hand drill, a color sprayer, etc. The TUI sits on top of a regular keyboard and mechanically translates the actions performed on it into keystrokes. Children can interact physically with the tools while the software challenges them with tasks, such as fixing malfunctions in virtual mechanical devices or model assembly.

Another example is Zowie's play sets [50] which use a low-cost tracking antenna and resonator tags [160] to track the identity and position of physical components that act as an interface to a PC game. In Redbeard's Pirate Quest (Figure 2.8), physical action figures can be placed in different parts of a pirate ship to support different interaction paradigms. For example, placing a figure near the steering wheel will enable the user to navigate the ship in the PC game, placing it near the cannon will shift the game into battle mode and placing a figure near the ship's telescope will switch the PC game's viewpoint to a magnified, "telescopic", one. Both Tonka Workshop and Zowie's play sets use simple spatial mapping between interface and task to help very young users achieve interaction levels that would be hard, or even impossible, to reach with the standard UI.

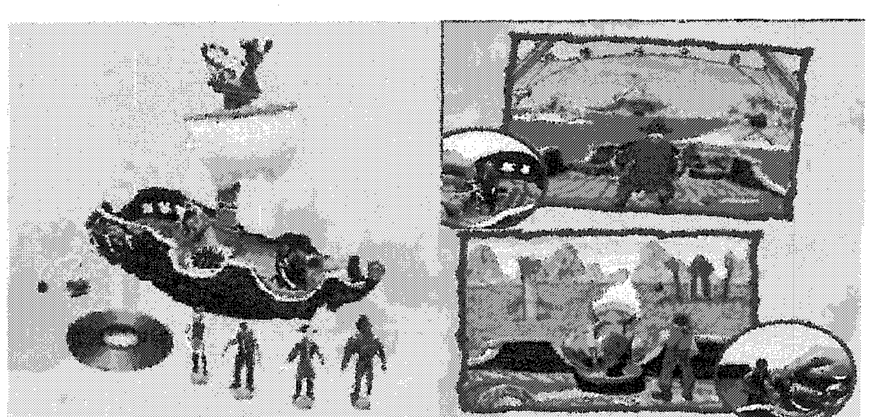

Figure 2.8 – Zowie's Redbeard's Pirate Quest

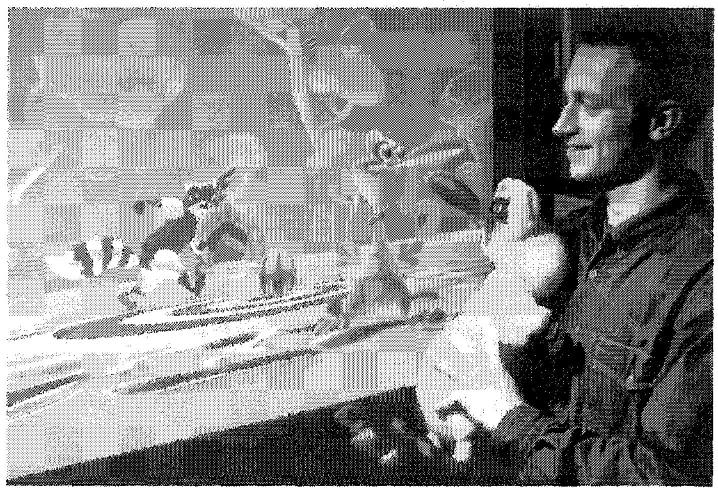

Figure 2.9 – The SWAMPED! interface

Plush toys afford natural behavior patterns to a child. The doll is inviting, huggable and supports pretend play [3]. The SWAMPED! Project [72] suggests an interesting way of controlling an animated character (Figure 2.9). The interface is a plush chicken toy affording an almost obvious spatial mapping from the physical doll to its virtual parallel. Control follows a voodoo doll metaphor: "Do to the doll what you would like the virtual character to do" [72]. The doll is embedded with sensors that sample its physical position. A "behavioral brain" is used to interpret the sensors output, for example, wobbling the physical doll back and forth makes it walk in the virtual world.

The SAGE project [164] stuffed rabbit storyteller agent (Figure 2.10), enables a child to interact with a digital storyteller. The SAGE rabbit can be a mediator to storytelling software running on a PC, or can act as a storyteller through a hidden speaker behind the rabbit's mouth. Physically changing the rabbit's hat transforms the storyteller's personality (for example, a Yarmulke for Hasidic stories, a Yin-Yang hat for a Taoist storyteller). The rabbit is also capable of rotating its

Figure 2.10 – The SAGE project's storyteller rabbit "Rabbi"

shoulders in order to change gaze direction, or move its ears and light its pupils to show attention and interest.

"Barney" (Figure 2.11) is a dinosaur "smart toy" designed for young learners [3,84]. Children treat him as a living being and can mimic and learn social behavior while interacting with him. Barney has several embedded sensors, including touch sensors in the feet and hands and a light sensor in his left eye for playing peek-a-boo. It also has some output mechanisms: motors at the head and arms and a loudspeaker. By being represented as a huggable, cute physical entity, "Barney" can achieve much better interaction with its users than a plastic keyboard.

Figure 2.11 - Interactive Barney with a friend

A few attempts were made to suggest an “Emotional TUI”, for interpersonal, tangible communication (non verbal or visual) [16,20]. Brave et al. suggested a “What You Feel Is What I Feel” remote tangible interaction paradigm, following the WYSIWIS – What You See Is What I See – paradigm, which is an underlying theme in computer supported cooperative work [16] (note that this is based on the classic WYSIWYG—What You See Is What You Get—paradigm emerging from Charles Simonyi’s Xerox PARC 1970s Bravo editor work [17]). Their implementation, inTouch (Figure 2.12), presented as a “tangible Phone”, is a set of two force feedback rollers connected over distance. Users who touch or roll their inTouch rollers feel their distant counterparts movements and thus, presence. LumiTouch (Figure 2.13) [20] uses a pair of distant picture frames as tangible interfaces between remote loved ones. When a user is present in front of her frame, or touches it, her presence or touch will be manifested as different glows of the distant picture frame.

Clearly “Huggable Interfaces” spatially map a friendly plush toy to an application. It is less clear how well the current emotional interfaces

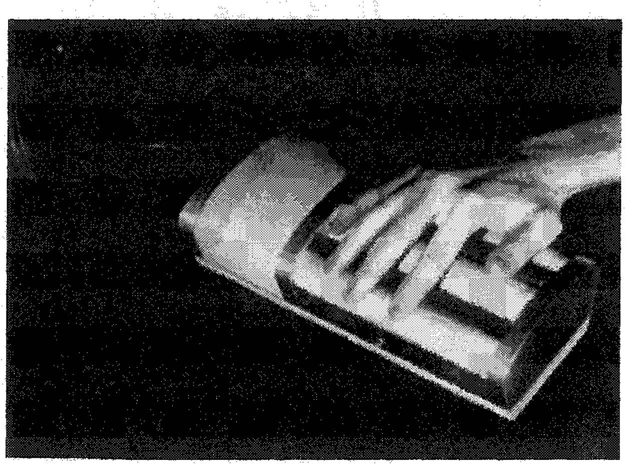

Figure 2.12 - inTouch


Figure 2.13 - LumiTouch

embody and convey emotions by touching a frame, or a roller. However, we believe that the path these set is very promising. The use of touch to express emotions to a remote person, or even the use of touch to input emotions to a digital computer, can reveal human sensations that will be hard or impossible to unveil using the standard UI.

2.7 Waldo Interfaces

Robert Heinlein's character Waldo [64] invented a series of remote manipulators, WALDOs, that enabled him to cope with his severe muscular weakness. The WALDOs echoed the shape and structure of the devices being controlled, whether huge cranes or nanomanipulators, affording intuitive control across scale. In many ways current telepresence and telerobotics research is attempting to implement practical WALDOs.

Two such WALDOs were developed for animators, enabling easy manipulation of virtual 3D models into a desired pose during keyframing.

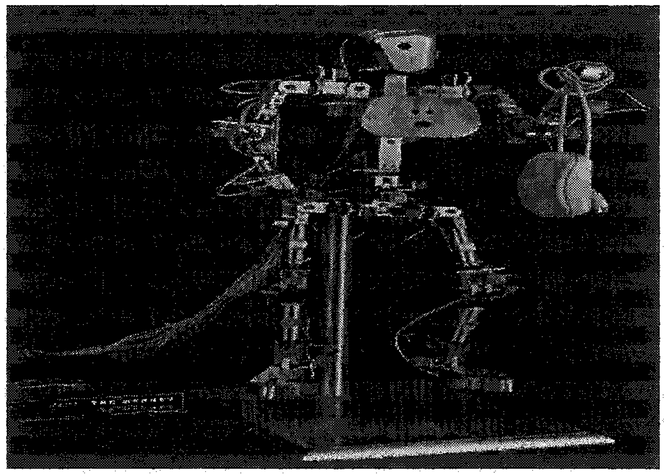

Figure 2.14 - The Monkeys interface, victorious over the common mouse

Both Dinosaur Input [34,40] (used for the film Jurassic Park) and Monkeys [38] (Figure 2.14) implant sensors in each joint of a mechanical skeleton to measure the joint angle. Since the topology of the skeleton is known *a priori*, a full representation of the skeleton's state can be sampled by determining the sensors' state. Dinosaur Input and Monkeys support an almost perfect spatial mapping between the physical interface and the virtual character's position and motion. Controlling virtual characters using the standard UI is dramatically more difficult and non-intuitive. However, the user's interaction with the Dinosaur Input and Monkeys interfaces is strongly restricted by the given and limited number and topology of the skeleton joints and their DOF. This limitation is a benefit when controlling certain types of characters but a weakness if the user wants to alter the characters' structure.

2.8 Tabletop TUI Paradigm

2.8.1 Overview

A number of researchers have explored TUIs based on a 2D surface (horizontal or vertical) that acts as the interaction domain and several physical objects that serve as interaction mediators. In this section we will attempt to present what we believe are the most important milestones of tabletop TUI research. As frequently revisited interface paradigm, these TUIs led to the emergence of some exciting technologies and ideas, including some allowing unusually strong I/O unification.

2.8.2 Early Examples

In RUGAMS, or "Real Reality" (Figure 2.15), tangible interaction is based on tracking the user's hand position and gestures, using tethered electromagnetic trackers and Data Gloves, rather than on tracking the physical objects that are acting as interfaces [18,135]. Based on known

Figure 2.15 – The RUGAMS, 'Real Reality' System

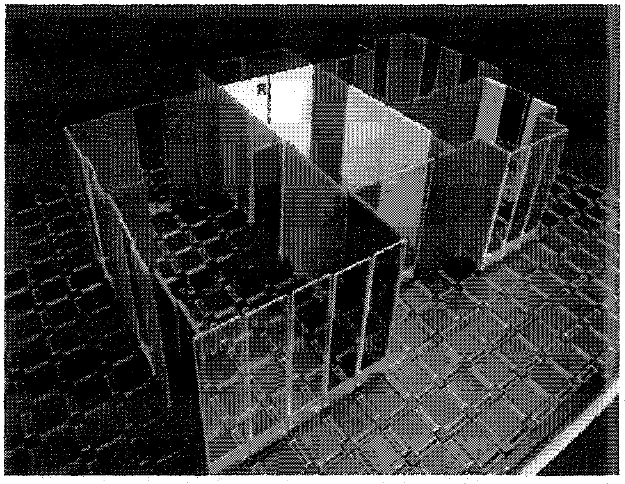

Figure 2.16 - The Segal Model

initial objects' positions, the system can track the location of the object being handled at any given time. Some of the reported applications of the system are tangible design of simple conveyor systems and an event-based logistical simulation of layout and material flow in a plant. "Real Reality" does not augment the interface by projecting graphics scenes on top of the interaction surface, limiting I/O unification.

As part of his pioneering "Machine Readable Models" (see also Section 2.11), the Segal model was built by John Frazer and his colleagues in the early 1980's to support the work of architect Walter Segal [52,53]. Segal had developed a technique allowing individuals to build their own homes but had found that the users encountered difficulties when it came to designing their homes by themselves. The TUI that Frazer developed enabled simple floor plan prototyping for users without any knowledge or experience with either computers or architecture. The interaction was mediated through the use of plastic panels, representing walls and windows, and a number of small scale wooden models of home appliances and furniture (Figure 2.16). The objects were connected physically and electronically to a grid of

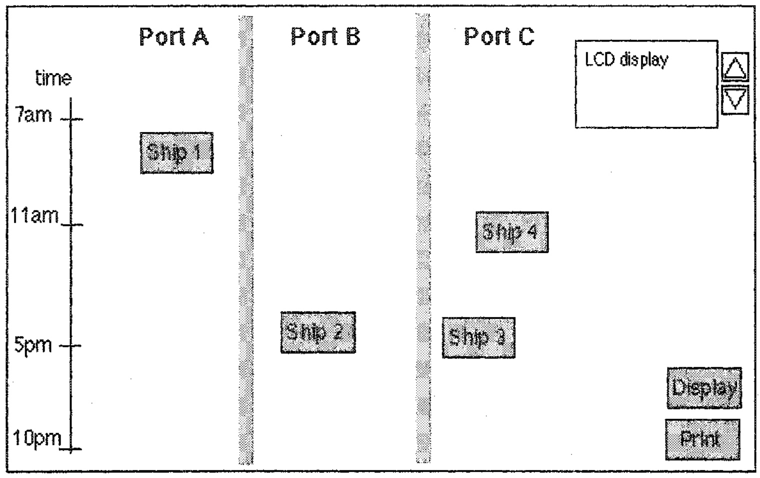

Figure 2.17 - The Lego Wall

connectors covering the surface of the model. The Segal model used a screen to display 3D wire-frame graphics of the resulting house model and did not augment the interaction surface with digital information, again limiting I/O unification.

Another early example, the Lego wall (Figure 2.17) was essentially a vertical tabletop TUI used for scheduling ships [39,40]. The Lego Wall was a set of blocks, each with an internal electronic ID, and a wall-panel that consisted of a 2D grid of connectors. Each block represented a different ship or simple actions (like "Print"). The wall-panel represented a table with time as rows and port names as columns. The model enabled users to follow and schedule shipping progress and access the ships' cargo information. The Lego Wall engaged human spatial awareness and skills, enabling users to simultaneously manipulate, sort and organize several physical objects tangibly, a very natural human approach to handling information items (a closely related recent effort is Senseboard, see Sections 2.8.4 & 2.12.5).

Fitzmaurice et al.'s Bricks [39,40] was part of the authors pioneering graspable user interfaces (essentially an earlier synonym for

TUIs) effort. Bricks were based on "physical handles" that were used to interact with virtual objects. The basic Bricks system employed generic, cube-shaped "physical handles". Bricks included three main prototypes: Bricks for Drawing, Bricks for Curve Editing and FlipBricks.

Bricks For Drawing consisted of a projection on top of a tabletop surface, called the Active Desk, and two physical "blocks" or "Bricks" tracked by two electromagnetic, tethered trackers. Bricks for Drawing supported a simple drawing application, GraspDraw, enabling the user to create, move, rotate and scale simple 2D shapes using the two Bricks (Figure 2.18). Similar themes were used in "Bricks for Curve Editing", using two blocks to adjust the contour and shape of a curve. A simple shape manipulation program as well as a menuing program called Flipbricks were implemented with non-tethered tabletop technology (Figure 2.19). Fitzmaurice et al. compared Bricks to the standard UI's mouse (see Section 2.12.3 for a discussion of this experiment). They argued that while the mouse is a "graspable device", Bricks, as an instance of graspable user interfaces, is a "graspable function" [39,40]. A

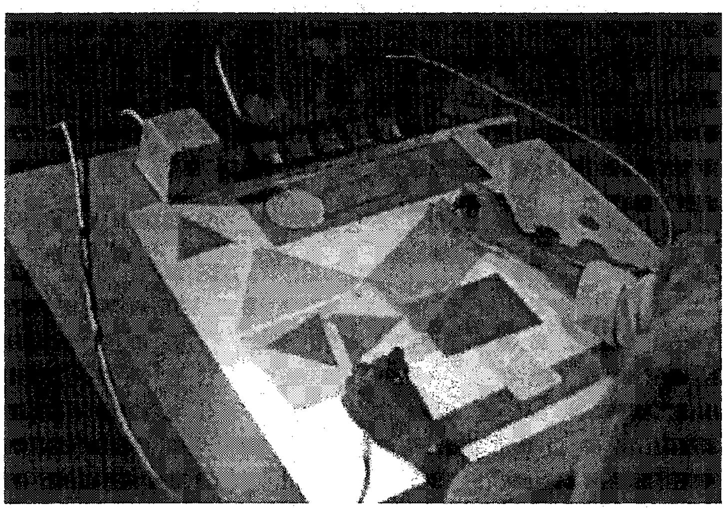

Figure 2.18 - 'GraspDraw'

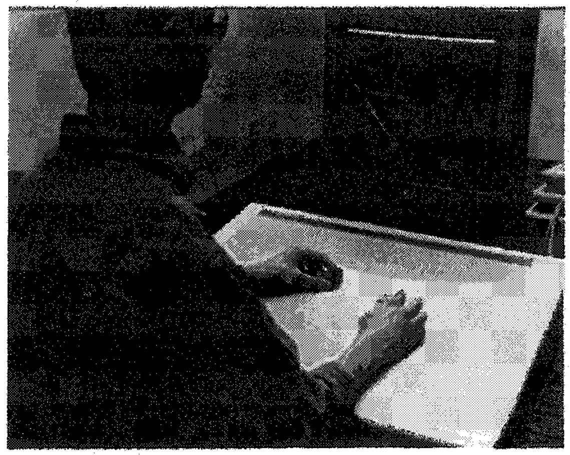

Figure 2.19 - Bricks shape manipulation application

"graspable device" is used in three phases: (1) acquiring the physical device (in the real world), (2) acquiring the logical device (in the virtual world), (3) manipulating the logical device. A "graspable function" on the other hand requires only two phases: (1) acquire the physical device and (2) manipulate the logical device. Fitzmaurice et al. also highlighted the benefits of improving affordances when they compared generically shaped Brick controllers and a specialized Brick based controller that had physical resemblance to the virtual shapes being manipulated by it.

Using Section 2.2's terminology we can argue that a "graspable function" is simply a result of adequate I/O unification. The benefits of affordances are given when a more intuitive spatial mapping is maintained between interface and task. Bricks was one of the first systems that attempted to implement these principles, and measure their benefits. However, Bricks allowed for only very simple and limited interaction and didn't enable detailed spatial shape and structure input.

BUILD-IT [42-48] is a more elaborate TUI that used physical blocks as interaction handles (Figure 2.20), tracked using a vision-based technique. The system also supports a direct projection of virtual scenes on top of the flat interaction medium, as well as on a vertical screen. The blocks are used for selection, movement and activation of virtual objects. The virtual objects, which can be much more complex than the simple blocks, are projected on top of the blocks, supporting I/O unification. BUILD-IT supports two (or more) handed interaction, and was designed for collaborative group work. Like Real Reality, BUILD-IT was initially used for prototyping of a plant floor plan, and more recently as a test bed for tangible scene navigation methods [44]. We will mention BUILD-IT again in Section 2.12.6 where we discuss a user study empirically comparing it to other interaction paradigms.

2.8.3 Ishii's I/O Bulb Systems

Hiroshi Ishii, one of the most prolific TUI pioneers, frequently used the tabletop TUI paradigm. The vision technology underlying many of his TUIs was what Ishii called the I/O bulb (Figure 2.21) [165]. The I/O bulb

Figure 2.20 – BUILD-IT

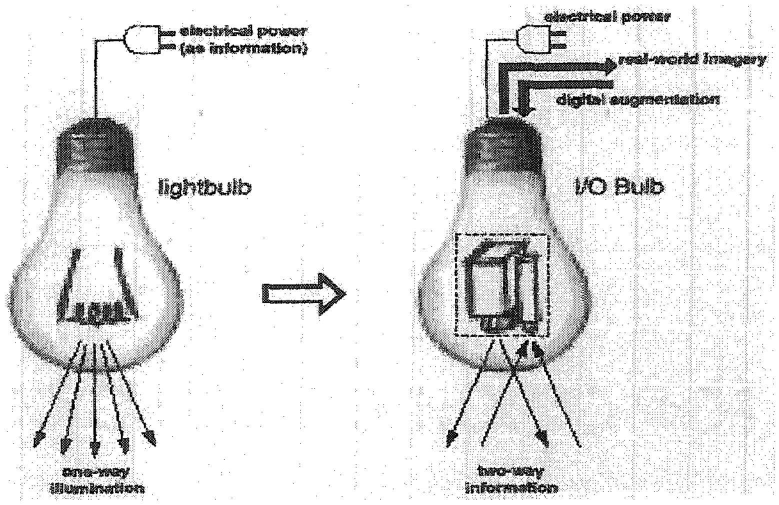

Figure 2.21 – The I/O bulb

could wirelessly track objects on top of a tabletop surface while simultaneously projecting detailed scenes onto the surface. The I/O bulb was the infrastructure for a number of applications, including an interface to a digital map in metaDESK [69,161], an optical design prototyping tool in “Illuminating Light” (Figure 2.22) [16,165] and an urban planning tool in URP (Figure 2.23) [166]. In all of these applications the user handles several tangible physical objects on the surface as means of controlling entities and actions in the virtual world. The surface is constantly augmented by digital feedback that is being

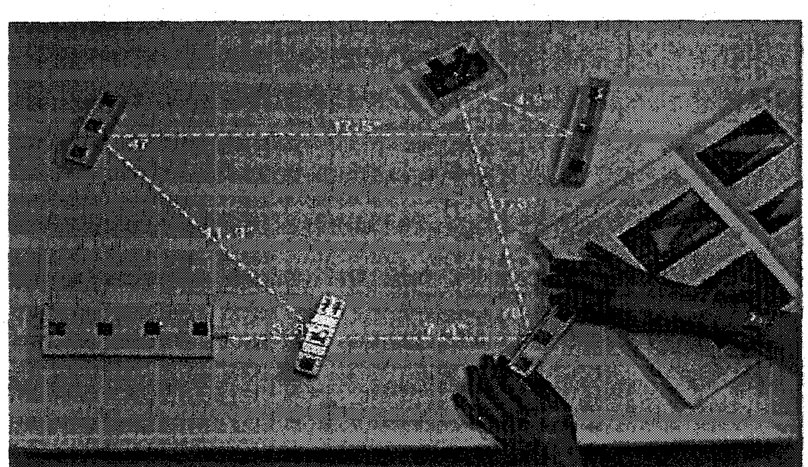
Figure 2.22 - Illuminating Light

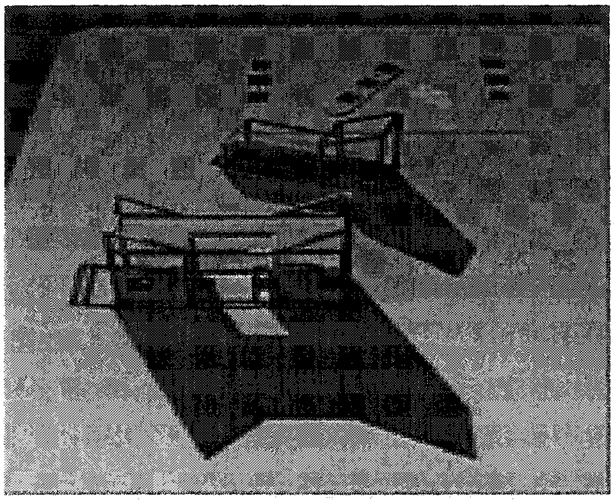

Figure 2.23 - URP

projected on top of it. For example, in URP the user can place a few physical models of buildings and roads on top of the surface, URP then creates feedback by projecting virtual building shadows. The user can use a physical clock model in order to change the virtual sun position, and accordingly the superimposed shadows.

The I/O bulb infrastructure proved to be a flexible and far-reaching TUI paradigm. A well designed TUI like URP supported spatial mapping of interface to task, I/O unification, and trial-and-error actions.

2.8.4 Other Tabletop TUIs

Several tabletop TUIs were based on Radio Frequency ID (RFID) tags that enable tracking of objects to which they are attached. Jacob et al.'s Senseboard (Figure 2.24) [71] uses tagged rectangular, magnetized pucks that can be attached to a vertical surface, being augmented by digital projection. Senseboard was used for an information organization application, a task that is still commonly performed by spatial placement of physical objects (for example, pieces of paper). Some of the pucks

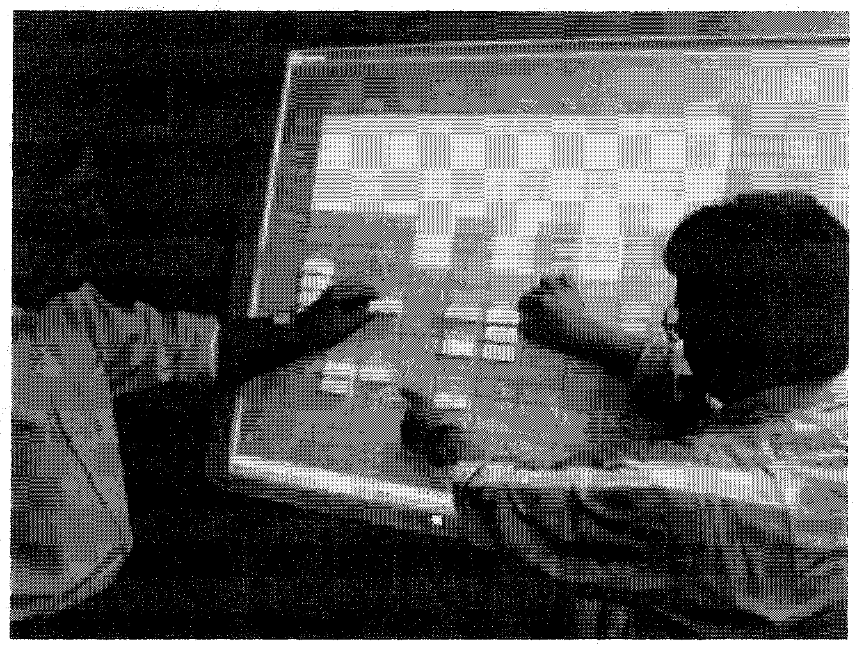

Figure 2.24 – Senseboard

represented "Data Objects", standing for a certain item of the information being organized. Senseboard augmented the "Data Objects" pucks by digitally projecting the heading of the information they contain onto them. "Commands" pucks, differently shaped, act as operators on "Data Objects" pucks, allowing grouping/ungrouping, detailed viewing, automatic constraint testing, save/restore, etc. Senseboard was designed to capture an existing, tangible, non-automatic, 2D-spatial approach to data organization. It offers a straightforward spatial mapping of the physical data organizational tool to the organization task, and enhances the existing manual approach with the benefits of automation. It also supports trial-and-error actions and I/O unification. Its limitation to the 2D spatial space is extremely well suited for its task. The Senseboard effort included an interesting small-scale user study that compared it to other interfaces, including a paper-based interface and the standard UI. We will touch on this comparative study in Section 2.12.5.

Figure 2.25 – DataTiles

DataTiles (Figure 2.25) uses RFID to track transparent panels that are placed on a flat touch-sensitive display screen [121,160]. The display augments the panels, creating a visual illusion of the information being projected from the panels (rather than through them). The interaction is deliberately based on the standard UI using the stylus and the panels as WIMP components. The panels could be held and moved on the display, and could be physically associated with each other to create simple modular construction which could represent a simple "physical language", symbolizing functions or behavior. By using a hybrid combination of TUI and non-TUI elements, avoiding task-specific spatial mapping, the authors argue that they can scale the interface to many applications, unlike more special purpose TUIs [121].

Tangible query interfaces (Figure 2.26) is an attempt to use a tabletop TUI for database queries [160]. The TUI is based (mostly) on RFID tracking and augmentation of the interaction surface using digital projection. A physical set of tools represents and controls the database

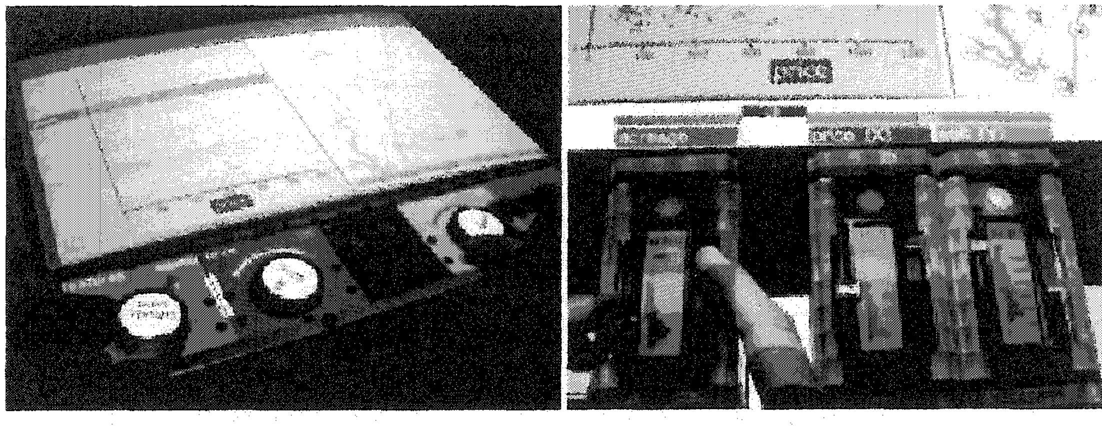

Figure 2.26 – Tangible query interfaces

parameters. The user can place these tools on a rack in order to express a database query. The results of the query are projected on the interaction surface and the user can edit the query interactively while examining the outcomes. Although tangible query interfaces spatially map physical objects to database queries this TUI raises an important question. Can spatial mapping between physical objects and a task that might not be spatial ever be intuitive? We will further discuss this interesting TUI, and empirical experimental results comparing it to non-TUIs, in Section 2.12.7.

Another variant of tabletop TUIs is based on electromagnetic wireless tracking using a multiple-stylus tablet, with the stylus components implanted in the TUI's physical objects. FlipBricks [39,40] (see Section 2.8.2, and Figure 2.19) is one example. A more recent work, Sensetable [110] supports a large, projection augmented, interaction surface and tracks multiple puck-like objects (Figure 2.27). Sensetable was able to solve some of the technical drawbacks of the other tabletop tracking techniques: it supports orientation tracking, a large number of tracked objects, does not limit the position of these objects to cells, and purported to be relatively reliable. Some of the applications suggested as

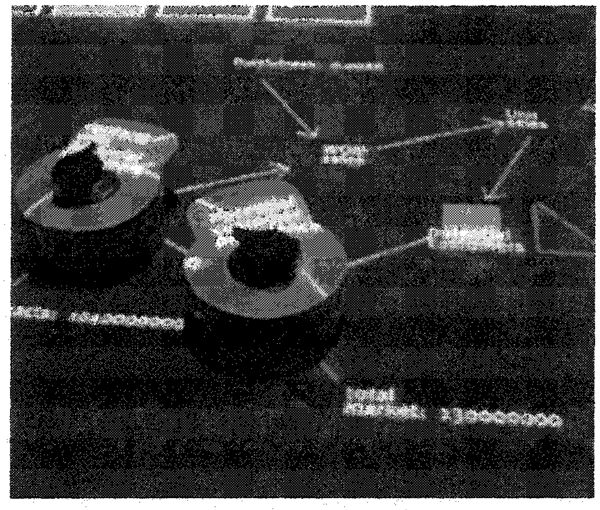

Figure 2.27 - Sensetable

being suitable for Sensetable were a chemistry educational tool [110], a musical synthesizer [111] and IP network simulation [81].

2.8.5 I/O Unification in Tabletop TUIs

Many of the tabletop TUIs mentioned earlier satisfy the I/O unification heuristic by digital projection on top of the interaction surface. Although the visual illusion of the projected information being part of the physical entities can be impressive, it is clearly not a strong physical binding and can be broken down completely by lack of tangible (haptic) display, shadows from the user's hand, or latency (delays in system response). Pangaro et al.'s Actuated Workbench [108] is an attempt to extend this coupling to the tangible haptic domain and, to more closely unify input and output. The Actuated Workbench is a tabletop TUI based on a set of optically tracked magnetic pucks that can be manipulated on a flat surface which covers an 8×8 grid of electromagnets (Figure 2.28). The pucks that can be moved freely by the user can also be moved by the TUI, using the electromagnetic grid (the

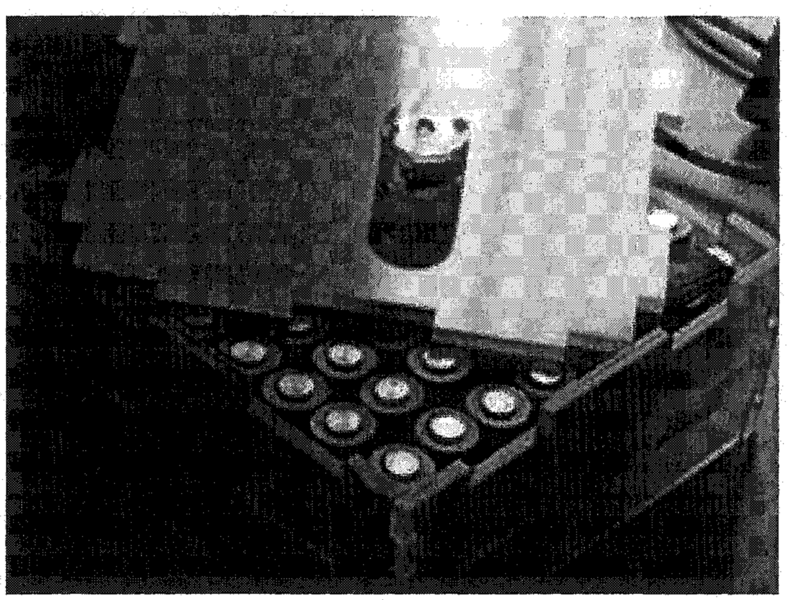

Figure 2.28 – The Actuated Workbench

authors also report that the electromagnetic grid can flip the pucks at will). Although the Actuated Workbench has not been employed yet for a specific application but is rather presented as a novel TUI technique, it opens the way for a new set of exciting tabletop applications, bringing TUIs closer to the scenario we presented in Section 2.1. The authors detail the following ideas for future applications of this technology [108]: “physical retrieval of past user actions, physical teaching and guidance tools, mechanical simulation, remote collaboration, physical management of information and entertainment ”.

2.8.6 Conclusion

The tabletop TUI family reveals a few of the weaknesses and promises of this emerging technology. Firstly, the large number of applications covered by Tabletop TUIs is notable. Examples range from planning (RUGAMS, BUILD-IT), design (the Segal Model, Illuminating Light, URP) scheduling (the Lego Wall, Senseboard) and drawing (Bricks).

Generally, tabletop TUIs can be designed to replicate and automate numerous surface oriented activities. Some of these TUIs offer very intuitive spatial mapping to the task (see for example Senseboard and URP), while others require high levels of abstraction (for example Tangible Query Interfaces or Sensetable's IP network simulation application). All of the tabletop TUIs support trial-and-error actions as they all offer multiple and persistent access points to the application's spatial state.

Secondly, tabletop TUIs have been used extensively as a natural test bed for new technological ideas and interaction paradigms. While the early generation of tabletop TUIs did not support I/O unification (for example, the Segal Model, RUGAMS, the Lego Wall) the second generation provided partial, visual I/O unification (for example, URP, Sensetable, BUILD-IT) and the latest example of the Actuated Workbench is the strongest TUI (not just tabletop) example known to us of I/O unification.

It is interesting to note that although tabletop TUIs have existed for more than a decade they have not been successful in penetrating any applied domain. While simpler TUIs successfully targeted commercial toys (see for example Music Blocks, Tonka Workshop or Interactive Barney) tabletop TUIs seemingly targeted more complex, "serious" tasks, but failed to deliver commercially. Should this failure be attributed to the current low level of expressiveness these TUIs offer? Should it be attributed to them not targeting the toy market? Are they simply too expensive for the tasks targeted? Or perhaps they have failed to find an applied domain that will benefit greatly from their functionality without requiring a higher level of expressiveness? These questions will likely be answered by developments in the coming years.

2.9 TUIs for Input of Vectors, Curves and Surfaces

Several TUIs were designed to support intuitive input of vectors, curves, surfaces and volumes. We see them as part of an effort to reach what is arguably the "Holy Grail" of spatial TUI research – "Digital Clay", a metaphor for media that would enable the user to mold complex physical shapes and at the same time would be sampled continuously into the digital domain.

A practical and simple approach to vector input was suggested in HandSCAPE [87], a TUI based on a measuring tape (Figure 2.29). While the user operates the HandSCAPE measuring tape, the measured magnitude and direction values are digitally sampled and transferred directly to a host computer. HandSCAPE affords straightforward and intuitive spatial mapping between the physical measuring tape and the application. An application of HandSCAPE, GeoSCAPE, is intended for

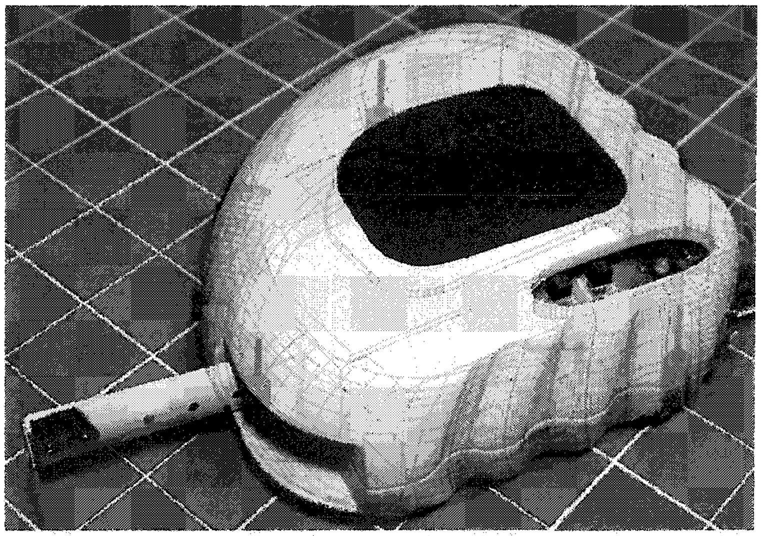

Figure 2.29 - HandSCAPE

on-site archeological excavation applications [86]. The GeoSCAPE approach is to directly sample the user measurements as vectors with absolute location, magnitude and direction so these can in turn be accumulated and eventually integrated into a 3D visualization of the archeological site.

ShapeTape (Figure 2.30) is a flexible strip that can sense its bend, twist, position and orientation [11]. ShapeTape is targeted at 3D modeling, helping the user to interact more naturally with curve primitives rather than working with a non-intuitive, mathematical spline-coefficients representation of the curve. The physical qualities of ShapeTape can be altered by attaching it to a constraint frame (for example, steel springs). However, ShapeTape doesn't allow tangible interaction with more than a single curve at a time. The physical characteristics of ShapeTape offer intuitive spatial mapping to a curve-editing task. In a very recent effort however a ShapeTape based interface was used as a 3D modeling tool, supporting curves input, and through these curves, 3D surfaces and solids modeling [62].

The haptic lens (Figure 2.31) supports interaction with a rubber

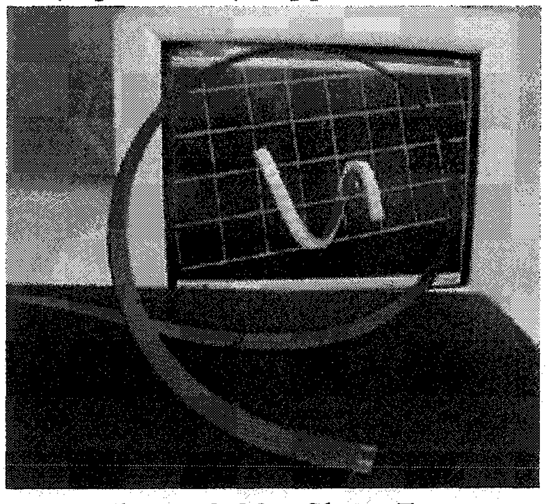

Figure 2.30 - ShapeTape

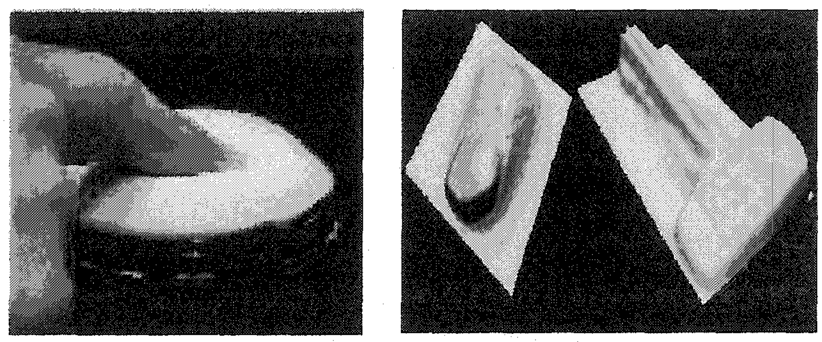

Figure 2.31 - The Haptic Lens

surface using fingers or objects [41,148]. The haptic lens samples its surface in real time, creating a digital height-field representation. However, this field is limited to a small surface and thus cannot express many shapes. The device uses a light absorbing material inside a membrane. An evenly distributed light source illuminates the material, while a camera images the light attenuation. Areas of the lens that are being pressed by the user will appear lighter in the image since the light passes through a smaller portion of the light absorbing material and undergoes less attenuation. Hence a height map of the surface can be extracted from the image. The haptic lens uses physical objects as molds; all that is required from the user is to push an object into the lens, affording intuitive spatial mapping to its application.

Another surface capture tool was an essential component of Illuminating Clay [115]. In Illuminating Clay the I/O bulb 2D interaction surface optical sampling paradigm (see Section 2.6.3) was enhanced with 3D surface optical capturing capabilities. Much like the I/O bulb, Illuminating Clay augments the interface by projection (Figure 2.32). However, the tool also uses a ceiling-mounted laser scanner to interactively sample the 3D topography, or height-field, of a clay model which is being manipulated by the user. The tool was designed to be an

Figure 2.32 – Illuminating Clay

interface for “real-time computational analysis of landscape models” for domains such as land design and engineering [115]. The clay model, which mediates the interaction, represents a topographic landscape (see also the variant of Illuminating Clay—SandScape—that uses sand, not clay [170]). The TUI generates and projects data such as slopes, shadows, solar radiation, water flow and land erosion back to the clay model [115]. Illuminating Clay offers an intuitive spatial mapping between the clay model and the 3D topography. Moreover, the clay, being a consistent 3D physical media that does not reset its spatial state to an initial condition when not touched, also supports trial-and-error actions.

3D surface manipulation was suggested in DO-IT [101]. The TUI is a deformable cube, enabling the user to reshape it simply by using two hands (Figure 2.33). The cube has a skeleton made out of conductive foam. The foam’s resistance is sampled continuously. Reshaping the cube causes changes in the resistance of the skeleton, enabling approximation of the surface deformation. The cube can be used to

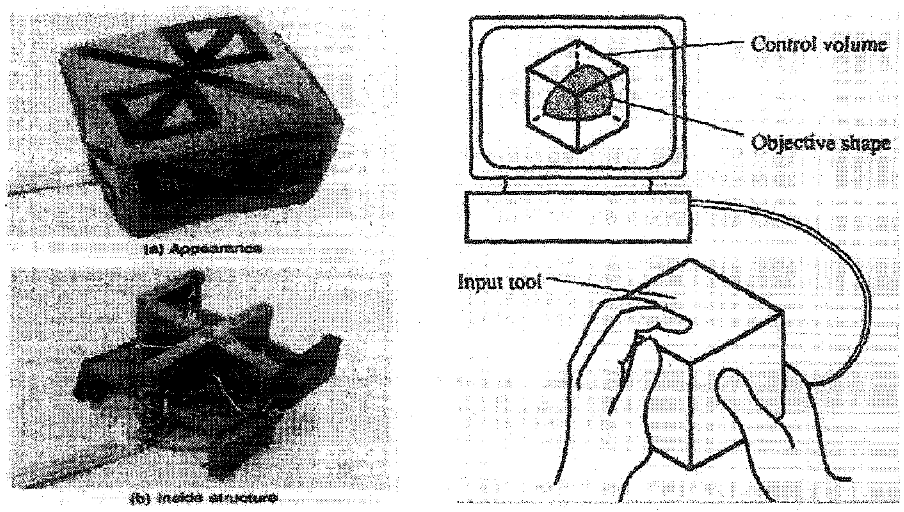

Figure 2.33 – DO-IT

roughly reshape simple virtual objects surfaces (not just cubes) by mapping the cube deformation onto the virtual object. The spatial mapping between the physical DO-IT cube and the virtual object being handled could be more or less intuitive, depending on the virtual shape that is being manipulated (for example, manipulating a virtual pyramid with the physical cube would require some abstraction from the user).

2.10 TUIs for Topology Input

Several TUIs suggested the use of a tangible, editable network of physical objects as a tool for topology input. These TUIs do not explicitly support shape input (except for simple 2D layouts), but still represent more obviously spatial TUIs since they offer intuitive physical means to edit the topology or structure of virtual spaces. AlgoBlock (Figure 2.34) was used as a computer programming education tool [155]. The tool consists of a set of physical blocks that can be assembled together by the user in order to build a computer program. Each of the blocks has an electronic ID and can be connected to a neighboring block on a surface.

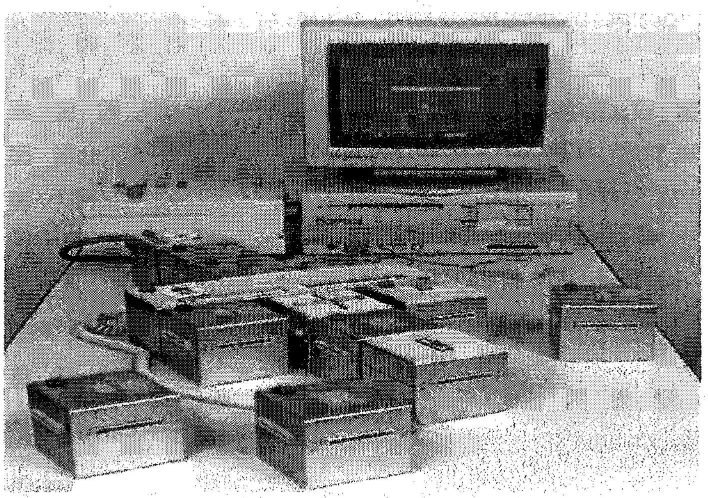

Figure 2.34 - AlgoBlock

Each block represents a Logo-like command and a program is compiled by interconnecting a sequence of several blocks to a host computer.

The Triangles TUI [59-61] consists of a set of plastic triangle-shape units (Figure 2.35). Each Triangle contains an embedded micro controller, a unique digital ID, and a picture drawn on it (for example, a drawing of Cinderella). The Triangles can be connected to each other with magnetic connectors, forming a network. This network is normally planar, but can be pseudo-3D by forming small pyramid-like shapes. Simple messages are transferred among the Triangles and to a host computer that continually samples the network topology. Changes in the TUI topology can result in various actions in the task domain. For

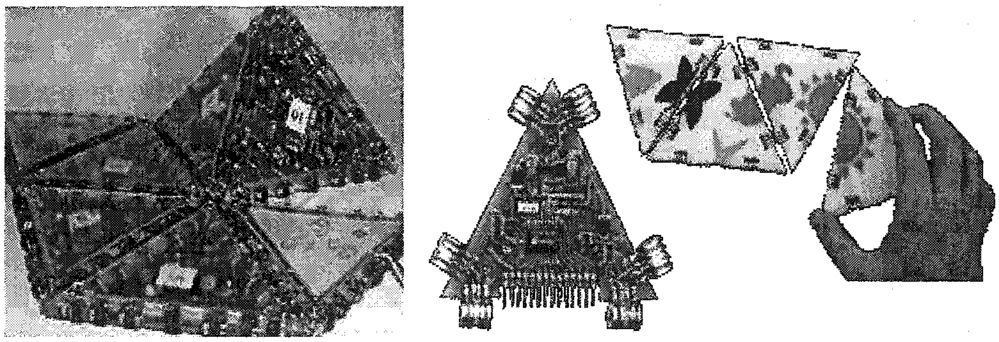

Figure 2.35 - Triangles

example, in a storytelling application, adding a "Cinderella mother Triangle" to a "Cinderella Triangle" will cause the mother to shout at Cinderella in the PC-based application.

Triangles are intended primarily for the input of connectivity and topology, rather than geometry and shape, and as far as we know are currently a tool looking for applications [59]. Proposed future research directions are nonlinear storytelling (for education applications and toys) and group coordination (tangible scheduling, workflow design, etc.) [59].

Triangles and AlgoBlock are multipurpose TUIs for 2D-topology input. Their basic shape hardly supports any specific spatial mapping between the physical tool and the application. This generic design can allow them to be spatially mapped to completely different 2D applications. At the same time their design also limits the quality of the resulting spatial mapping between them and the application. Triangles and AlgoBlock do encourage trial-and-error actions but do not offer strong I/O unification.

2.11 Input of 3D Shape, Volume or Structure

TUIs directed solely for detailed 3D spatial shape and structure input have been developed over the last 20 years by several research groups in different disciplines (from architecture to mechanical engineering), mostly unaware of each other's work. Some of the first such working interfaces were the pioneering UIs developed by John Frazer and his colleagues as early as 1980 [51-55] (Figures 2.36 and 2.37). The UIs enabled feedback supporting the architect. It is important to note that Robert Aish contemporaneously published and developed tools along similar lines [1,2].

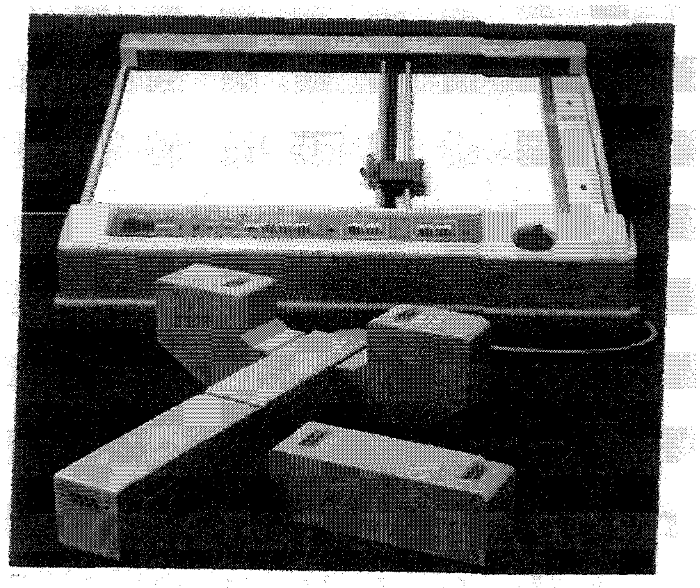

Figure 2.36 – Flexible Intelligent Modeling System

Figure 2.36 and 2.37 present the “Flexible Intelligent Modeling System” and the “Three Dimensional Intelligent Modeling System”, respectively. Both could be used for 3D shape input by simply attaching and detaching the models’ blocks. A host computer could then sample the physical model that was assembled. Unfortunately, until recently the HCI community did not refer to these early UIs and many of them were

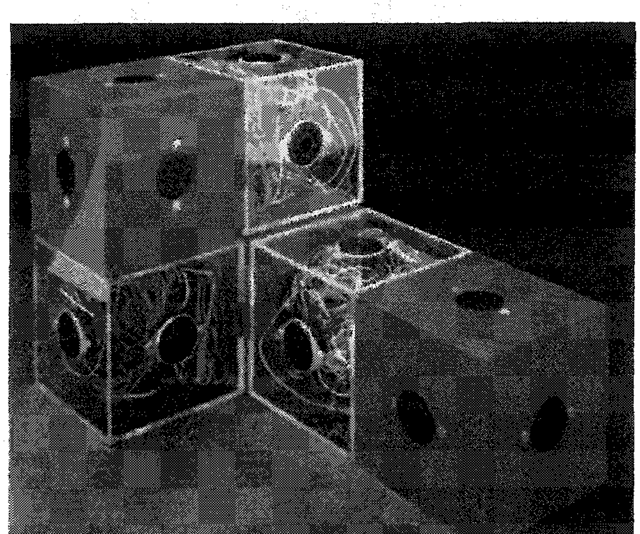

Figure 2.37 – Three-Dimensional Intelligent Modeling System

“reinvented” by various research groups.

The Geometry Defining Processors (GDP) presented in 1989 [5] were designed as a support tool for engineering system definition and analysis. GDP are a set of blocks, each containing a CPU and a means of communicating with its neighbors and with a host computer. The blocks are modular, enabling the user to attach and detach them. The user was expected to interactively input the geometry of a dynamic system or of an engineering problem by building it with the blocks. After the geometry of the problem was tangibly described, the blocks attempted to solve the problem they describe by communicating between each other and with the host computer. The reported application of the system was a thermal optimization design of a simple fin. The system did not address user interaction issues and as far as we know was not used in other applications.

Recent efforts enable a user to define, in an easy-to-use manner, a fairly large-scale 3D geometry by simply building it with blocks. Each block contains some processing power and can communicate with its neighbors and through them with a host computer. By browsing this 3D network the host can extract its topology and render a virtual representation of the physical spatial structure.

MERL’s blocks [6,7] are a stackable, large (10x5x2.5 cm), Lego-like TUI for 3D structure input. The blocks were demonstrated in a very detailed, 560 block structural design (Figure 2.38). The blocks however cannot be sampled interactively. Instead, MERL’s blocks are assembled off-line and, only when the construction is done, digitally sampled in a relatively slow process (as an extreme example, sampling the 560 block

Figure 2.38 – MERL's blocks

structure in Figure 2.38 took 53 minutes). MERL's blocks were demonstrated in toy and PC game level prototyping applications.

Mark Yim's (Xerox PARC) Digital Clay project [175] is another 3D TUI aimed at 3D structure or shape input. It is a derivative of Yim's work on deformable robots (see Section 2.13). Digital Clay (Figure 2.39) is a set of dodecahedron (12 faces) modules each the size of a large marble (~2.5cm in diameter). Each module contains a small processor and each face has a unique ID. Each module can be connected to the other

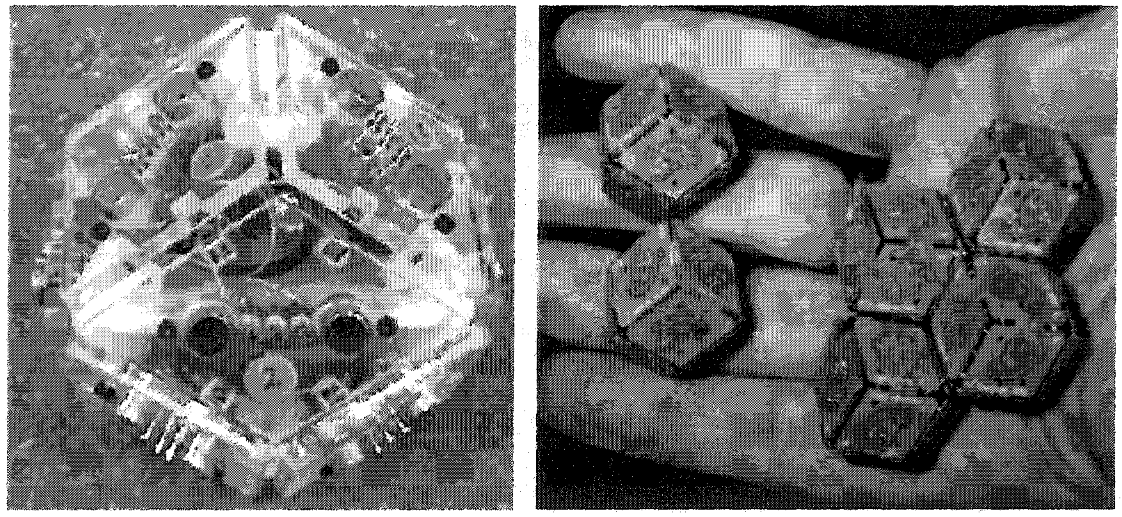

Figure 2.39 – Digital Clay

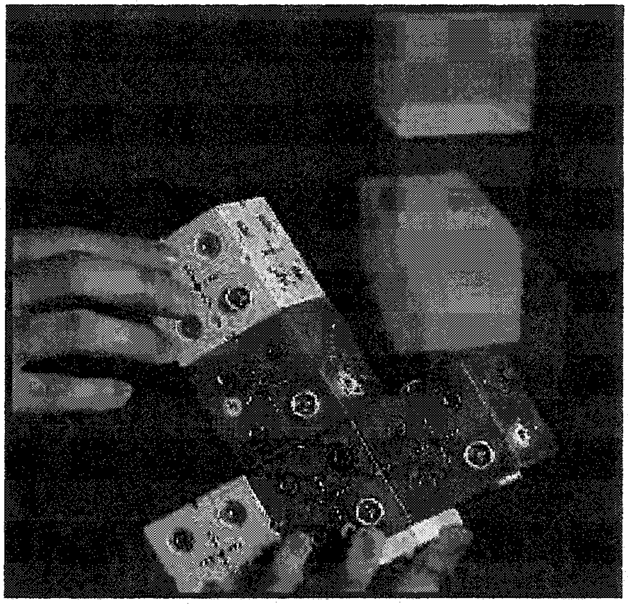

Figure 2.40 – ActiveCube

modules using magnets, establishing a network through electrical connections. The tool still seems to be in early prototyping phases and we are not aware of it being used for a specific application, or of its performance in practice.

Yoshifumi Kitamura and his colleagues at the Kishino Lab at Osaka University developed ActiveCube, a TUI for 3D shape input and multimodal interaction [70,78-80]. ActiveCube (Figure 2.40) supports 3D construction using a set of plastic cubes (5 cm/edge) that can be attached to, and detached from, other cubes through any of their six faces. The connection is established by male-female connectors (similar to cloth snaps) forming both an electrical network topology and a physical shape. A host PC samples the network and the structure online, registering connection and disconnection events in real time.

ActiveCube are also equipped with a variety of input and output devices: ultrasonic, optical (visible and IR), tactile, gyroscopic and temperature sensors, and light, audio, motor and vibration actuators.

Through these, ActiveCube can support multimodal interaction—in a way, forming an editable, interactive 3D robot [70,78-80]. We believe that currently ActiveCube is the best existing example of a 3D spatial TUI for structural input.

How close to ideal are these TUIs designed for 3D spatial shape, volume and structure input? Let us examine them in light of our tangible interaction heuristics (Section 2.2). Most significantly, these novel tools are the first TUIs presented that afford intuitive spatial mapping between their physical shape and a detailed 3D input task. The user can physically build a complex 3D structure and edit its physical spatial shape, while the virtual domain samples and follows. It seems that these tools should make designing virtual 3D shapes and structures much easier even for inexperienced users. On the other hand, these tools have physical characteristics that are quite limiting. The current tools are fairly big (the working prototypes have modules with physical dimensions roughly the size of the larger Lego-Primo blocks) and generally uniformly shaped. Even though there are very few limitations to the size or complexity of the structure the user can build with these TUIs, the physical vocabulary that they offer is very limiting.

These TUIs can offer very good I/O unification. In many applications the TUI serves as both input and display, unifying action and perception space and reflecting the state of the design task by its physical existence. An exception to this ideal is device error, where the TUI's physical state is not registered correctly causing the physical state to disagree with the application state. A more severe exception is the case where the physical shape is a proxy for a more detailed, complex, smooth or abstract virtual shape which is displayed externally from the TUI. Due to the current limited expressiveness of 3D spatial shape and structure input TUIs, this can present a real obstacle, breaking I/O unification whenever the user is required to view the impact of her work on a separate display.

The 3D spatial shape, space and structure input interfaces we presented naturally support, much like any Lego blocks, trial-and-error actions and activity. The user can perform a series of straightforward steps in order to form a physical structure representing her cognitive design goal. At any point, the user can also use the TUI to edit the physical structure in a trial-and-error manner, not following any linear pattern, by simply attaching or detaching any number of blocks.

We can conclude that these new TUIs bring us closer than ever before to the spatial TUI goal of a natural interface between human and computers in 3D spatial shape, space and structure input tasks. The major drawback of these tools is their currently limited spatial vocabulary which might limit their application to a confined set of tasks.

2.12 Comparative Studies of TUIs

2.12.1 Empirical Evaluation of TUI vs. Standard UI

Are TUIs actually "better" than the standard UI? If so, when? Given the huge diversity of TUIs and their applications, and the spatial strengths of the standard UI (see Section 2.4), this question could be difficult to answer. In this section we will summarize the few attempts made to experimentally make the TUI/standard UI comparison.

As we shall see, TUI effectiveness varies by application. In fact, in several studies, the TUI was less useful than the other tested UI.

Nevertheless, we believe that all these studies point to the important need to carefully tailor a TUI to its task. TUIs can stand or fall according to the quality of spatial mapping they offer between the physical tool and the task.

2.12.2 Rauterberg's "Go-bang" Study

In 1995 Rauterberg et al. [118] performed a large scale evaluation of the benefits and drawbacks of a TUI. They compared a DigitalDesk-like TUI (see Section 2.5.2) with non-TUIs in a game task, with the computer as an opponent. The game was "Go-bang", won by placing five pieces in a row. The four interfaces examined used the mouse, a touch screen, a command line interface (typing coordinates for moves) and a DigitalDesk-like TUI, which enabled the player to use physical chips while augmenting the game board with projected computer chips. Hundreds of participants were included in the study that took place at a large computer fair in Switzerland. The experiments were informal, with very little control. According to a questionnaire answered by the 304 participants the touch screen was the easiest interface for the game, followed by the mouse, the TUI and the command line interface. Interestingly, significant correlation was found between age and the usability score for the TUI: the older the player was the more likely that she or he gave a high usability score to the TUI. In actual trials, during 3,801 automatically observed, completed and non-drawn games, players using the command line UI won more often than the players that used the mouse. Players using the TUI won more often than the players that used the touch screen. Unfortunately, unbalanced experimental control did not allow direct comparison of the mouse and command line UIs to the tangible and touch screen UIs. The authors attempted to rebalance the results post hoc, reaching an overall conclusion that the TUI "guarantees a significant increase in the user's performance" [118].

This study is unique because of the unusually large number of participants. However, we believe its results should be taken with a grain of salt due to its informality and unbalanced control. We further believe that evaluating the interface usability by players' winning chances is arguable at best.

2.12.3 Fitzmaurice's Bricks Study

Fitzmaurice et al.'s pioneering Bricks project (Section 2.8.2) involved an attempt to empirically demonstrate that its proposed TUI paradigms offer advantages over their standard UI parallels [39,40].

In two experiments 10 users were asked to perform simple tasks. In the first experiment they were asked to manipulate a 2D virtual shape by translating, rotating and scaling it in order to match a static virtual prototype shape. In the second they were asked to acquire and track a 2D virtual object. The three interfaces tested with these two tasks were a standard UI-like pointing stylus, two generically shaped Brick controllers and a specialized Brick based controller with a physical resemblance to the virtual shape (improving the spatial mapping quality). While the user had to move between translation, scaling and rotation modes in order to match or track the shape with the pointing device, the tangible Brick controllers (whether generic or specialized) afforded simple manipulation of the virtual shape by physically translating, rotating or "scaling" the Bricks.

Both experiments showed a significant improvement of performance when using the Bricks controllers (generic or specialized) as interfaces instead of the pointing device. In the shape manipulation task, for reasons that the authors describe as "task-related" [40], there were no significant benefits to using the specialized Brick interface over the generic Brick. In the acquisition and tracking task, however, the specialized Bricks interface significantly outperformed the generic Bricks interface.

2.12.4 Rasa vs. Paper

The Rasa paper-based TUI for a military command post (see Section 2.5.2) was tested in a small user study that included an empirical comparison between the paper based TUI and the non-augmented sticky notes [96]. The study participants included nine officers in a pilot study and six officers in the final phase. The study design forced participants to evaluate both interfaces during one interaction session, without revealing to them the underlying comparative experimental motivation. Participants were asked to start a command post session using Rasa, and then witnessed a (controlled) power failure that forced them to use the paper interface solely until the power, and Rasa with it, were turned on again.

Based on questionnaire and interview feedback, users reported Rasa to be an easy interface to use in comparison to the non-augmented paper, and overall preferred it to the classic sticky note interface. Based on objective comparisons of action times when using Rasa and when using the paper-only interface the authors report that there is no evidence that using Rasa benefits or penalizes user interaction time budget, excluding time needed for Rasa's error repairs. They also note that this is not much of a surprise as Rasa is closely based on the physical sticky notes as interaction mediators. A separate early study hints that searching for spatial information on the map, a heavily time-consuming manual process using the paper only interface, could be dramatically improved using Rasa [96].

2.12.5 Senseboard vs. Paper and the Standard UI

Jacob, et al.'s Senseboard (see Section 2.8.4) was tested in an interesting user study that included comparisons to a traditional paper system, a pen-based version of the standard UI, and for experimental purposes, a more primitive version of the Senseboard (discussed below) [71]. Thirteen participants were asked to perform an information organization task—scheduling group work—to satisfy a set of constraints.

The "Reduced-Senseboard Condition" [71] removed the automatic notification when scheduling constraints were violated, leaving only the points in which the TUI is inferior to the physical paper. Through this setting the authors hoped to gain an empirical measure of the TUIs imperfection in replicating the real world, flaws that, arguably, are not an inherent part of the TUI concept and can be corrected as TUI technology improves.

Both the "Reduced–Senseboard Condition" and the Senseboard suffered from several TUI limitations. The tangible pucks introduce latency between participant actions and resulting display, lower display resolution than paper, and disappearance of the details on the puck when it is taken off the Senseboard. Based on questionnaire feedback the authors report a weak user preference for the TUI, and significant dislike of the paper system. In an objective assessment of tasks time-to-completion measure, the authors report results that were only marginally significant statistically. These results show that the Senseboard was fastest, trailed by the paper-only, standard UI and then the reduced TUI conditions [71].

2.12.6 Morten Fjeld's BUILD-IT Study

Morten Fjeld and his colleagues compared BUILD-IT (see Section 2.8.2) to other traditional interfaces [48]. The study task, inspired by the BUILD-IT's plan layout application, was a related spatial 3D constructional design task, involving pointing a laser towards a target.

The study compared four systems:

(1) BUILD-IT (an essentially 2D TUI) application in which the user could manipulate a virtual laser beam by interacting with a BUILD-IT brick on the surface.

(2) A 3D physical small-scale replica of the positioning problem, including physical blocks and an operative laser.

(3) A 2D physical abstraction of the positioning problem, based on cardboard blocks and a laser representation (a metal ruler).

(4) A mathematical tool, basically a text definition of the problem, a calculator, piece of paper and pencil.

A system using the standard UI was not included because the authors felt that the highly 3D nature of the problem made it completely impractical for novices without lengthy training. The mathematical tool was totally outperformed by the other systems in a pilot study and was therefore excluded from the final study. The final experiment included 30 participants, each attempting to solve the positioning problem with only one of the three interfaces. The results show that both the physical tool and the BUILD-IT TUI outperformed the 2D cardboard tool. User satisfaction with the physical tool was higher than with BUILD-IT, although the time-to-completion measures were comparable. The authors summarize the experiment results by saying that "...the cognitive support offered by a ...TUI comes close to the physical world..." [48]. While not arguing this conclusion, we point out the inherent inferiority of the 2D tabletop TUI when it comes to 3D oriented tasks. We believe that an automatic 3D spatial TUI version of the tested 3D physical tool would have been a better fit for this task than BUILD-IT.

2.12.7 Ullmer's Tangible Query Interfaces vs. the Standard UI

Ullmer's work on tangible query interfaces (see Section 2.8.4) included a well-designed user study, testing the usability of the TUI and comparing it to a standard UI [160]. The study was conducted using a "home finder" real estate application. Participants were required to search a database of properties looking for a specific type of house. Each participant was given criteria for each search task and could define her own query using four independent search parameters (for example, look for houses with "minimal taxes, maximal size, near location B" [160]). The application continuously estimated the participants' performance by measuring the resemblance between their query results (that is, properties found) and the real answer to the given criteria. As soon as the resemblance reached a given threshold the task was declared as completed and the system recorded the task's time-to-completion. The user could express queries using the standard UI by moving WIMP-based virtual sliders, or by manipulating the TUI's physical set of sliders.

The study included 16 participants. It showed that the TUI was slower than the standard UI, and that user satisfaction was about the same for both the standard UI and the TUI. The author suggests several possible explanations for the TUI's poor performance, including the need to physically set the TUI's sliders before defining a query, and that two of tasks were harder to perform than the rest. Excluding the TUI's setup time and the outlier tasks the TUI performs as well as or even better than the standard UI [160]. The work presents several insights to these outcomes, including a reminder that the faster technology is not always the better, with reference to examples of the early standard UI era, when some tests showed the new standard UI, WIMP-based interfaces were slower than their text-based counterparts.

We believe tangible query interfaces illustrate the possible limitations of TUIs. We argue that although spatially organizing pieces of information in a physical way can be extremely natural on various tasks (as was successfully demonstrated, for example, in Senseboard and Rasa), this is not always true. In this case, the spatial mapping between a physical tool and the abstract notion of a database query might not be very intuitive. Do we naturally think in a spatial manner when we try to search databases with Boolean queries? Although this question should remain open and reexamined thoroughly, we believe that a TUI must above all satisfy the essential requirement of intuitive spatial mapping between the physical tool and the task. We argue that when this requirement is not satisfied, the standard UI, or perhaps even a textual interface, might prove to be superior to the TUI.

2.12.8 Conclusion

Let us reexamine our initial question—are TUIs "better" than the standard UI? Based on the examples in this chapter, the answer is at best inconclusive, at worst negative. All the TUIs covered in this section were successfully implemented, and trial experiments confirmed that they were useful and supported the task for which they were designed. However, they did not always perform significantly better than their non-TUI parallels, and in some cases even performed worse.

These are the very first attempts to measure TUI benefits over other interfaces. Most of the evaluations we reviewed were based on small statistical samples and some suffered from flaws in the experimental design. We further believe that some of these studies simply choose to hunt for the wrong prey, or to avoid the right one. For example, we agree with Fjeld's hunch that comparing BUILD-IT to a standard UI based 3D CAD tool will probably be an easy win for the TUI (Section 2.12.6), however, we would argue that such a study should still be performed. On the other hand, we suspect that tangible query interfaces (Section 2.12.7) might be an example for a TUI that, in its task setting, is not more effective than the standard UI. To conclude, we believe that the correct answer to our question is more toward the inconclusive than the negative. In any case, we see a clear lesson here: TUIs should not be dogmatically perceived as being "better" than the standard UI.

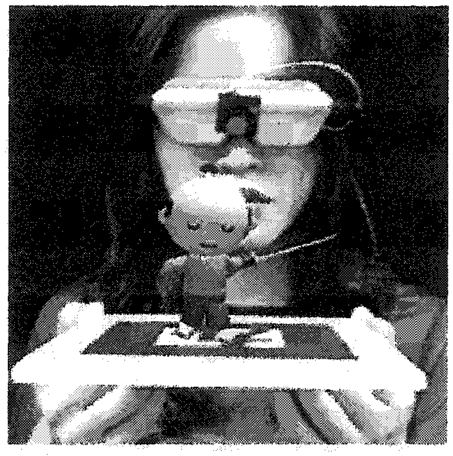

Figure 2.41 – MagicBook

2.13 Related Technologies and Systems

As we mentioned earlier, many other current research efforts are closely related to the TUI domain, but remain out of the scope of this literature review. For example, Billinghurst, Kato, Poupyrev and their colleagues developed Mixed Reality (MR), or Tangible Augmented Reality interfaces that blur the borders between the physical and the virtual interaction domains [13,73,116]. The MagicBook [13] (Figure 2.41) enables participants to manipulate a physical book that can turn into an augmented reality object (displaying virtual objects on top of its physical pages), or transform into a virtual sphere, allowing the user to "enter it" and be completely immersed in a virtual scene. Using similar technology, Tiles [116] targeted an aircraft instrument panel design task. Unlike tabletop TUIs, Tangible Augmented Reality visually augments its physical interaction mediators without restricting the interaction area to a surface. However, its tethered headset restricts the user and the physical mediators it employs are currently limited to flat, relatively large tiles, due to tracking requirements.

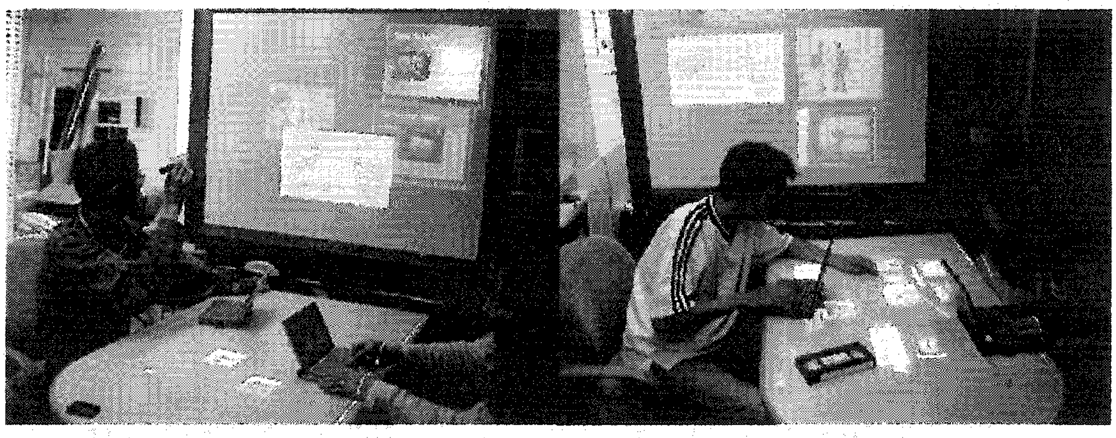

Figure 2.42 – 'Augmented Surfaces'

The idea of using practical everyday objects as interfaces is not limited to paper (see Section 2.5.2). The "The Office of the Future" [117] presented an interesting technological vision for the mixture of physical everyday office entities with digital virtual information. "Augmented Surfaces" [120] presented the use of notebook computers and other common office objects (for example, documents folder, video cassettes) as interfaces to digital information on a "spatially continuous workspace", projected on physical surfaces in the office (Figure 2.42).

Current work on reconfigurable modular (or deformable) robots also has close relation to TUIs. Mark Yim's PolyBot and Telecube (Figure 2.43 [176,177]) are modular entities that are "aware" of their structure and can deform and change it by autonomously attaching and detaching their modules, in order to support different functionalities (in Section 2.11 we mentioned Yim's Xerox PARC work on Digital Clay, a TUI derivative of his deformable robots). These entities might revolutionize detailed 3D spatial shape and structure input. In the long run, we can see this technology dramatically enhancing TUIs clarity of state (see

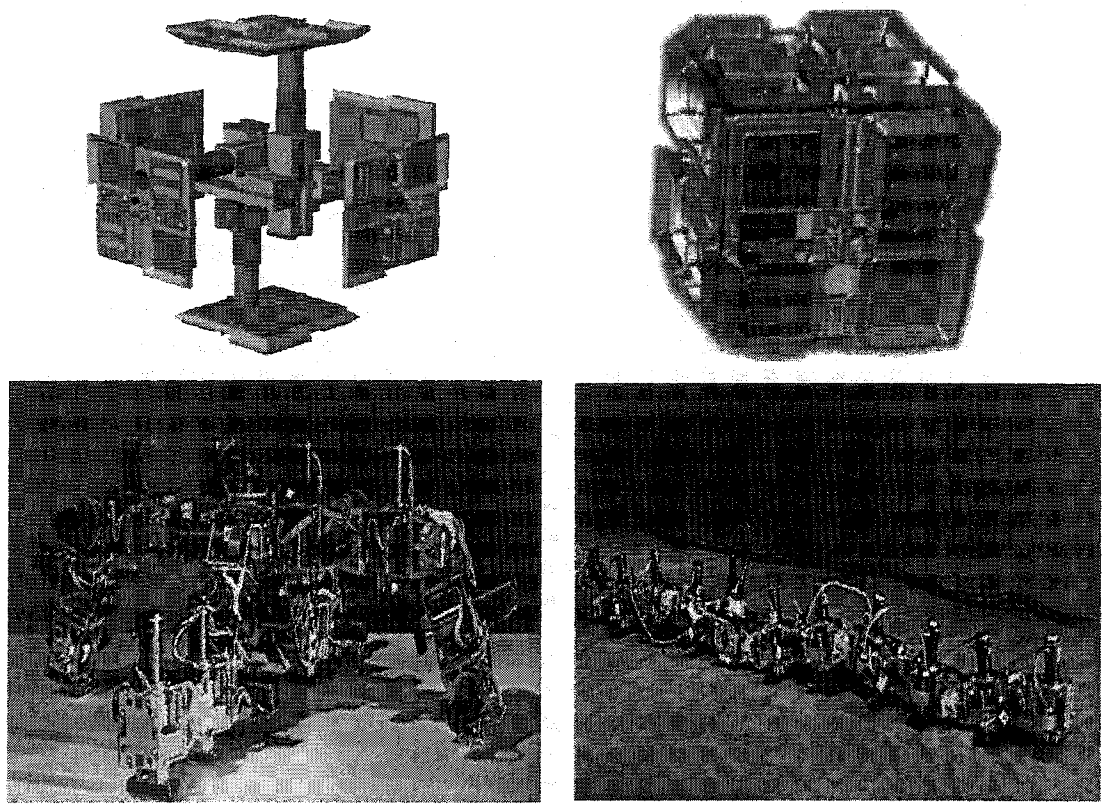

Figure 2.43 - Telecube (top) and PolyBot (bottom)

Section 2.2.3). A 3D TUI with deformable qualities could physically mirror the application state by essentially reshaping itself.

The NuMesh project [171] suggests a 3D modular network of high performance computers plugged into a 3D-lattice topology in a Lego-like modularity. Given such modularity and self-awareness combined with high computing power at each module, one can envision [134] a smart TUI which is actually a powerful 3D modular parallel computer (somewhat similar themes were implemented in the GDP project, see Section 2.11). Much like in a finite-element analysis, the user can build a physical representation of a spatial mathematical or physical problem by assembling the modular-computing units. Given the correct rules and initial conditions, the physical computer network that the user built represents the problem spatially and can simulate a solution, or "solve itself", by communicating between the modules.

Another related technology is 3D scanning [41]. 3D scanners use either active-vision (for example, projecting laser light onto the object) or passive-vision (for example, stereovision). Although such scanners enable the user to build a physical shape and then scan it directly into the digital domain, they all suffer from occlusions. Light cannot travel into non-transparent objects and so optical scanners are limited to scanning only the external surface of the 3D object and not its interior structure. Another problem relates to user interaction. Although constant scanning is possible (for example, 3D surface scanning in Illuminating Clay, Section 2.9), optical scanners are far from offering an ideal solution for interactive input of 3D shape or structure. Erroneous measurements caused by occlusions from the user's hands, along with the need to work within the field of view of the scanner (or constantly rotate the 3D object to scan different aspects of it) could make an interactive TUI employing such scanners cumbersome and probably impractical for many interactive applications.

2.14 What Lies Ahead and Closing Thoughts

Focusing on future trends in spatial TUIs, we are hoping to see growth in the application vocabulary they afford. We can foresee, in the short term, that TUIs like the ones presented in Section 2.11 will reach the size and expressive richness of regular Lego blocks. Such future tools will be extremely useful for early prototyping and design in many fields of engineering, science and art.

Another valuable technology for spatial TUIs might come from the growing availability of 3D printers. These printers can print relatively detailed and complex virtual objects into solid physical objects in a straightforward, automated manner, much like an inkjet printer, and are becoming more affordable. Such printers might bring new possibilities and new users to the early prototyping domain—users can build physical prototypes without the need of an orthodox workshop. Spatial TUIs might make the virtual design phase of this prototype cycle much easier, closing the gap between design and output, without requiring the use of complex CAD software. This also introduces the possibility of "beaming matter" or faxing 3D physical objects over digital communication lines, without ever needing to edit them in their digital form. Such a process could start by interacting with a TUI that automatically samples the interaction outcomes into digital form. The digital data could then be transferred to any other location and printed as a physical object using a 3D printer.

In the long run we can envision a complete fusion of the virtual and physical domains. Thinking about "active deformable" materials [84], we picture a system that would support input of 3D physical shapes and structures and at the same time facilitate a physical "display", that controls and deforms the physical medium. A straightforward application of such an I/O device would be a two way "digital clay"—a material that would: (1) enable detailed physical 3D sculpture-like input, directly sensing its shape in real-time; and (2) support physical output, interactively changing and deforming its 3D structure according to need. This would be almost ideal I/O unification.

Although such technology is not accessible at the moment, it might be within reach in the foreseeable future. Much effort is being put into micro and nanotechnologies. In the near future "chips" will integrate mechanical, optical and electronic components, all produced at the same micro or nano metric scale. It is not so far reaching to conceive micro- or nano-scale devices that could be assembled to a deformable physical volume, sense their location inside it and sample the volume's shape interactively. The same micro-devices could also output changes in shape by physically moving themselves or pushing their neighbors [84].

We believe that TUIs have an important role in tomorrow's computing world as additions to or, in some cases, replacement of the current standard UI. As we mentioned in Sections 2.3 and 2.4, the mouse has been widely accepted in the last 30 years of computing as a simple and successful means of utilizing users' tactile abilities. The standard UI is generic and offers specific affordances only rarely. With this generality, the standard UI has been successful in serving a wide variety of tasks, spatially mapped according to need, a quality that went hand-in-hand with the existing paradigm of personal computing.

Current computers and PCs with the standard UI are extremely generic devices [107], and so the mouse with its typical lack of affordances fits them well. This means that PCs generally remain complex devices, hard to master by the consumer. The PC's generality leaves many challenging applications and tasks poorly served. 3D spatial input, one of these tasks, is an obvious target for specialized TUIs that cannot offer the generality of the standard UI, but can outperform it by affording intuitive spatial mapping to a specific application. We argue that as today's generic PCs turn over their place to tomorrow's specialized information appliances [107], there will be a growing need for a variety of customized tangible interfaces that will complement, or replace, the standard UI. In order for future spatial TUIs to excel, they must be highly specialized, striving to adhere to the following design guidelines:

(1) Fully afford intuitive spatial mapping to specific application tasks. Exploit spatial abilities and mappings learned innately and early in life before those learned later.

(2) Further exploit users' 3D spatial abilities by supporting trial-and-error actions, and unifying input and output. This maximizes the usefulness of the mapping since it reproduces spatial settings in the real world.

(3) Explore the rich real-world vocabulary of physical objects, tools, and related spatial techniques as inspiration for novel TUI design.

We conclude with a condensed overview of the TUIs described in this chapter (Table 2.1). The TUIs are presented in the order in which they were discussed, and are evaluated according to the quality of the spatial mapping, I/O unification, and support of trial-and-error actions that they offer. These parameters are assessed using a subjective key (bad, fair, and good). The table also details whether or not the TUI affords 3D interaction, and some of the TUI's current applications. All of the judgments made here should be taken with a grain of salt, since we have not in person seen most of these systems in operation.

It is important to note how hard it is to technically unify input and output. Several TUIs choose augmentation of the action space with digital projection, which offers a visual illusion of merging action and

TUI (and associated chapter section)	Spatial Mapping	I/O Unification	Trial and Error	3D?	Current Applications
The Mouse (Standard UI) (2.4)	✓	—	✓	No	Generic, supports the WIMP interface
Simple Information containers (2.5.1)	✓	—	—	No	Media editing, toys
Paper TUIs (2.5.2)	✓+	✓	✓+	No	Information organization, air traffic control and military command post applications
Passive real-world interface props; Tonka Workshop; Zowie's play sets (2.6)	✓+	—	—	No	Neurosurgical visualization, toys
SWAMPED!, SAGE and Emotional Interfaces (2.6)	✓	—	—	No	Toys, tangible communication
Monkeys and Dinosaur Input (2.7)	✓+	—	✓	Yes	Animation keyframing
RUGAMS, The Segal Model, Bricks – FlipBricks (2.8.2)	✓+	—	✓+	No	Floor plan design, plant design, manipulation of virtual shapes
Lego Wall (2.8.2)	✓+	—	✓+	No	Information handling, ships scheduling
Bricks - GraspDraw, BUILD-IT, URP, Illuminating Light (2.8.2-2.8.3)	✓+	✓	✓+	No	Floor plan design, plant design, urban design, optical design
Senseboard (2.8.4)	✓+	✓	✓+	No	Information organization
DataTiles (2.8.4)	—	✓	✓	No	Generic, supports a hybrid of GUI and TUI interaction
Tangible Query Interfaces (2.8.4)	—	✓	✓+	No	Real estate database applications
Actuated Workbench (2.8.5)	✓	✓+	✓+	No	Generic, not demonstrated yet with specific application
HandSCAPE (2.9)	✓+	—	—	No	Measurement tape TUI
ShapeTape (2.9)	✓+	—	✓	No	Curves input
Haptic Lens (2.9)	✓	—	—	Yes	Surfaces input
Illuminating Clay (2.9)	✓+	✓	✓	Yes	Land design and engineering
DO-IT (2.9)	✓	—	—	Yes	Reshaping of virtual objects
AlgoBlock, Triangles (2.10)	✓	—	✓+	No	Educational toys
"Machine Readable Models", MERL's Blocks, Digital Clay, ActiveCube (2.11)	✓+	✓	✓+	Yes	Structure input & modeling

Table 2.1 – TUIs overview

Key: ✓+ *Good*
✓ *Fair*
— *Bad*

perception space, enhances clarity of state by altering the input device appearance according to the application state, and thus, to some extent, unifies input and output. As we discussed earlier (see Sections 2.8.3 and 2.8.5), though potentially very useful, this illusion is limited even in the visual perception realm, and does not unify input and output in the other sensory domains. We believe that a combination of visual augmentation with automated physical control over the TUI's action space could offer a novel and powerful TUI paradigm in the future. Currently we think that only SenseTable (Section 2.8.5) can claim to offer strong visual and tangible unification of input and output.

Given the hardships of technically implementing the input and output unification, it is interesting to note the ease of support of trial-and-error actions. Many TUIs support trial-and-error actions by simply allowing physical handling of pieces of paper, pucks on a surface, blocks in 3D space—all of which offer multiple points of interface access and tightly coupled mappings, enabling trial-and-error exploration of the task space.

We should emphasize that "3D Dimensionality?" is not a parameter in our assessment of the TUI's quality. As we mentioned earlier, a 2D TUI is by no means inferior to a 3D TUI (see for example Senseboard, Section 2.8.4, for a robust 2D TUI which is well designed for its 2D task).

Lastly, looking at the table's "Applications" column, it is interesting to note that while many TUIs do offer intuitive spatial mapping between physical objects and a specific application, some choose a more generic approach. Much like the shape of the mouse in the standard UI, many TUIs maintain a shape that is quite generic, for example the pucks in Senseboard and SenseTable or the modular components of Triangles, MERL's blocks and ActiveCube. These may be a first generation of generic TUIs that will support a variety of tasks, striking a useful balance between generality and specificity.

Chapter III
Cognitive Assessment

3.1 Introduction

3.1.1 Overview

When addressing the question of intelligence, Edwin Boring said in 1923 that "intelligence is what the tests test" [14,124]. Similarly, any assessment of human cognition is shaped, and limited, by the tools it employs. Examining the technological component of cognitive assessment techniques reveals a stagnant state. In 1997 Robert Sternberg highlighted that intelligence testing changed very little over the last century, making little use of the powerful technology presently available [125,127,150]. The canonized Wechsler Adult Intelligence Scale (WAIS) appeared in 1939 and has changed little since. Assessments that were computerized frequently followed classic paper-pencil based tests as prototypes, barely scraping technology's potential [125,150].

Many research endeavors are targeting this deficiency. Since cognitive assessment is all about attempting to have a glimpse of human cognition, state of the art HCI technology should have a dramatic impact on the field. VR, a far-reaching HCI paradigm, is already being exploited as a research test bed for a number of novel cognitive assessments (see for example [124-126,136]). We see TUIs playing an important role in pushing forward the field of cognitive assessment.

In this chapter we will briefly overview a few concepts of cognitive assessment, focusing on the issues underlying and justifying our research. This chapter is not a profound overview of the field of cognitive assessment, instead it merely supports the concepts and terminology used in our work and in this dissertation.

The chapter starts with a short discussion of cognitive and neuropsychological assessment (Section 3.1.2). It then briefly overviews the current efforts to automate cognitive assessments, emphasizing the emerging role of VR and the potential impact of TUIs (Section 3.1.3). Section 3.1 ends with a look at aging and age-related issues that are addressed in our experiments and in this dissertation (Section 3.1.4). Section 3.2 discusses cognitive mapping and its assessment; Section 3.3 discusses constructional ability and its assessment with a quick look at the mental rotation test.

3.1.2 Cognitive and Neuropsychological Assessment

Cognitive and neuropsychological assessments are scientific attempts to study cognition even though it cannot be approached and observed directly [99]. Examining cognition by directly monitoring physiological brain activity could fail to deliver clinical value unless matching behavior is also understood [124]. With this in mind, neuropsychology has evolved as *the science of evaluating and determining specific physiological brain activities by examining observable human behavior* [89,124,125,136]. Cognitive assessment also involves the measurement of human behavior, but is more scientifically oriented and less clinical than neuropsychological assessment [99].

Testing human behavior involves giving the participant an opportunity to "behave" and measuring it. A measurement tool should be reliable (yielding the same results consistently on different occasions) and valid (measure what it is supposed to measure) [124,152]. Obviously the measurement tool should be sensitive, safe and should offer the assessor full control over the data collection process [136].

Allowing the participant "to behave" involves the presentation of stimuli which trigger recordable reactions by the participant. Arguably, many classic, paper-pencil cognitive assessment tests offer very limited stimuli, little freedom to behave and low ecological validly (that is, little relevance to normal, everyday human behavior in the real world) [124].

3.1.3 Automation of Cognitive Assessment

Most major psychological paper-pencil tests have been automated or are expected to be automated in the near future [63]. Immediate advantages for this kind of automation are saving in professional's time: the computer tirelessly samples the participant actions and reliably stores, and refers to vast assessment knowledge, dramatically reducing the expertise requirements from the assessor. Other obvious advantages of automation are extremely high density of measurement, elimination of tester bias and potential improvement in test reliability. Computerized tests can also be sensitive to response latency, and enable questions tailored based on the examinee's past answers [63]. Automated assessment has also been criticized with concern focused on miscalibration of tests with respect to their written parallels and misuse of tests by unqualified examiners.

We share the view that this kind of straightforward automation portrays merely the tip of the iceberg for automation, and that much of the naysayers' arguments against automation are based on tradition rather than on scientific vision. Robert Sternberg suggested automation-supported "dynamic assessment", where tests targeting learning offer guided performance feedback to the participant [127,150]. Major efforts address the potential of VR for cognitive assessment.[3] VR-based cognitive assessment should afford all the obvious benefits of automation, particularly almost ideal assessment reliability [124,136].[4]

Albert Rizzo and his colleagues (and others) promote VR-based cognitive assessment as a breakthrough in the field of neuropsychological assessment, enhancing assessment validity and everyday relevance [123,124,126,136]. VR-based cognitive assessment can "objectively measure behavior in challenging but safe, ecologically valid environments, maintaining experimental control over stimulus delivery and measurement" [136]. VR-based cognitive assessment also introduces many new challenges. One largely unaddressed need is the analysis of huge number of measurements the automated tools extract ("drowning in data" [124]), compared to the simplistic measures of traditional assessment (commonly a single time-to-completion measure per task).

Examples of VR-based cognitive assessments are quite different from the classic paper-pencil tests. In the Virtual Classroom (Figure 3.1), which targets assessment of Attention Deficit Hyperactivity Disorder (ADHD) [123], the participant sits at a desk in a realistic 3D interactive

---

[3] Recently, the coupling between VR, cognitive assessment and psychology was tested on the flipside with a VR application measuring its realism (or its Presence: *the sense of being there*) by the psychological responses of its users [98].

[4] Cognitive assessment was presented as possibly being the "kinder killer application" for VR technology, a technology that in the past was driven mainly by military applications [129].

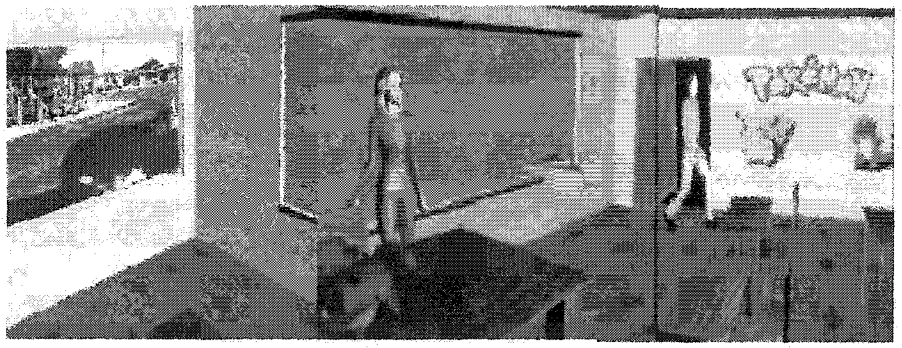

Figure 3.1 – The Virtual Classroom

simulated classroom. The participant attempts to follow the instructions of a virtual teacher while ignoring distractions such as cars passing and children in a nearby playground. Although affording a seemingly "real-life" setting, the Virtual Classroom maintains full control over assessment stimuli and measures.

We see an important role for TUIs in enhancing the ecological validity of assessments further. Spatial TUIs, earlier defined as *tangible user interfaces used to mediate interaction with shape, space and structure in the virtual domain*, can stimulate and measure numerous human behaviors that would be hard or impossible to measure without them. A straightforward example is the potential use of TUIs for automatic assessment of constructional ability (see Section 3.3 and Chapter 4). Less obvious is the use of TUIs as mediators to spatial human behaviors that require a more indirect mapping to the interface. Special care should be taken to maintain the task ecological validity when the spatial mapping between the human behavior and the TUI is less direct (for example, see the automatic assessment of cognitive mapping ability, Section 3.2 and Chapter 5).

3.1.4 "The Keepers of Culture"

Cognitive assessment can be used for numerous goals; for example, "head hunting" of technically gifted employees. We chose to concentrate on assessments that could potentially benefit the elderly population and their caregivers.[5]

The population in the developed world is rapidly aging. In North America alone millions of baby boomers (the 80,000,000 or so babies born between 1946-1964) are nearing retirement age. This dramatic demographic process was termed by Horace Deets the "aging revolution" [32]. The "aging revolution" poses many challenges for society, not all of them predictable. Gene Cohen would like to see the elderly regaining their traditional role as the glue that holds societies together, or "the keepers of culture" [22]. Deets, like many others, sees an extremely important role for technology in successful aging processes and advocates a convergence of the two (technological and aging) revolutions [32].

Technology can benefit the elderly community in numerous aspects of life. One such aspect is the assessment (and rehabilitation) of cognitive and functional skills. Normally, healthy elderly suffer from decline in some mental skills, without a matching loss in others [89]. Skills that are known to decline with age are those that require speed and involve active solutions of complex new tasks. General spatial ability

---

[5] Other than ethical reason, our rationales for addressing elderly-related issues were quite practical: 1) following our collaborators interests; 2) The relative accessibility of healthy elderly participants; and 3) The well-known sensitivity of a number of cognitive abilities to age.

is known to decline with age, affecting abilities such as mental spatial visualization, spatial memory, spatial relations and mental rotation (See Section 3.3) [66,89]. High-level abilities such as wayfinding and cognitive mapping can also decline with age (See Section 3.2) [76,156]. Decline in cognitive mapping ability in healthy elderly can vary and be affected by emotional aspects such as familiarity or attachment to known places [9].

The elderly population is more prone to suffer from various brain diseases and dementia [89]. Most notably is the high proportion of the elderly suffering from Alzheimer's disease (AD), the most prevalent type of progressive dementia [89].[6] AD is still very hard to assess precisely, especially in its early stages, and definitive diagnosis is possible only through autopsy [75,90]. AD affects several spatial skills, including mental rotation and constructional ability (see Section 3.3), and can severely damage high level spatial skills such as cognitive mapping and wayfinding in previously unknown environments [75,89-92,126] (see Section 3.2). Even in its early phases AD can dramatically hinder a person's ability to perform everyday activities that were previously well within her capabilities, like driving or finding her way in a new place [90-92]. This is true to such an extent that the missing person waiting period is waived for diagnosed dementia patients, who have died from exposure when they become lost and disoriented. Currently AD is diagnosed by a physician such as a neurologist or geriatrician relying on exclusion criteria and using specific diagnostic criteria [75,89-91].

---

[6] For example, in the 1990s about 8% of Canadians aged over 65 suffered from AD [75].

The reasoning we presented earlier (Section 3.1.3) in support of automating cognitive assessments becomes even more compelling when thinking of the elderly. VR-based assessments can dramatically enhance the ecological validity and present the elderly participant with an appealing and realistic environment that does not have the feel of a fearsome test on one hand nor a less-than-worthy childish puzzle on the other (see for example [90,122,126,136]). However, issues such as computer skills and simulation sickness should be carefully addressed when approaching elderly participants with VR-based cognitive assessments (see for example [93,156]).

We believe TUIs can be an extremely powerful means for approaching the elderly community with new technological assessment tools. A well-designed TUI can almost completely hide its technological components and leave the elderly participant with a physically manifested mental challenge based on a set of tangible objects spatially mapped to the task. Participants can perform the physical assessment task completely unaware of the automation controlling the stimuli and measuring their every move.

3.2 Measuring Cognitive Maps

3.2.1 Cognitive Maps

In his pioneering 1948 paper, "Cognitive Maps in Rats and Men" [158], Edward C. Tolman argues that rats, as well as humans, have a mental representation of the world—a cognitive map as he termed it. These cognitive maps hold detailed spatial information that the animals collect, integrate and use while they interact with their environment.

Tolman's cognitive maps can vary in their resolution and accuracy from a comprehensive map to a narrow strip according to parameters such as the level of training and the cognitive state and abilities of the animal. In a series of maze tests Tolman and his students proved that rats build a spatial cognitive map of the environment even in the lack of a stimulus such as food. Later, when they are presented with a stimulus, they use the cognitive map that they acquired in order to reach their goals faster. The cognitive map can be quite comprehensive and can give the rats a high level of spatial orientation. Rats that were trained in a maze found a short way to their food even when their regular path was blocked or when the maze's walls were taken away, obviously referring to a global sense of direction and knowledge of routes. These concepts have led to the modern psychological definition of a cognitive map: *an "overall mental image or representation of the space and layout of a setting"*, and in turn cognitive mapping is: *"the mental structuring process leading to the creation of a cognitive map"* [8].

It is important to distinguish between the psychological concepts of wayfinding and of cognitive maps. Wayfinding is the *"cognitive element of the overall process of reaching a destination* or *the cognitive element of navigation"* [31], and involves the following spatial problem solving components: *"decision making, decision executing and information processing"* [8]. Cognitive maps underlie the wayfinding process and enable making and executing decisions about the environment, connecting a decision execution plan to a spatial cognitive map [8]. Simply put, having a precise cognitive map of an environment is necessary but not sufficient for reaching a destination.

Understanding the qualities of cognitive maps and their creation is an extremely important concern in various applied fields, from urban planning to army training.

3.2.2 The LRS Model

The process of cognitive mapping is part of our everyday interaction with environments and environment representations. Our interaction can be a direct and active "physical" interaction, like walking through a neighborhood or diving around a coral reef, or an interaction mediated through a variety of indirect means (for example, a map, video clip, or a virtual reality walkthrough). Both physical and virtual environments can be valid means of acquiring a cognitive-mental representation since both are external to the learner [58].

The true nature of cognitive maps is not well understood; currently, the most widely accepted theory of cognitive mapping is the Landmark-Routes-Survey (LRS) model [31]. The model divides our environmental understanding into three hierarchical categories that can be integrated into a single comprehensive cognitive map [24,31,58]. Landmark (or declarative) knowledge is the simplest level of understanding an environment. Landmark knowledge contains a list of objects that exist in an environment. Landmarks will usually be dominant features of the perceived landscape; the person holding the landmark knowledge will be able to declare their presence and name them. Route (or procedural) knowledge uses the landmarks as markers and decision points, and adds topological procedural information connecting the landmarks by a series of paths and travel distances with specific orientation. Route knowledge includes a sequence of routes, distances, turning points and actions that should be taken between the landmarks. Performing this sequence will bring the traveler from landmark to landmark. Survey (or configurational) knowledge is based upon the integration of the landmark and route knowledge, representing the environment as a topographical-geometrical configuration of objects and routes in a fixed and more precise global coordinate system [24,31,58,157].

Different kinds of interaction with an environment lead to different amounts of knowledge in the three categories. Route knowledge can be achieved by egocentric sequential interaction with the environment (for example during navigation or orienteering). Survey knowledge can be based upon route knowledge but can also be acquired directly through map learning [8,157]. Acquiring detailed survey knowledge will allow better estimation of distances in the environment than estimations based on route knowledge. At the same time, survey knowledge based solely on a map will lead to worse navigational performance than knowledge acquired by actual navigation [8,31,157]. Gaining more experience in an environment usually transforms the cognitive map from route knowledge to more abstract survey knowledge [157]. Furthermore, different kinds of interaction with an environment can lead to completely different kinds of cognitive maps. For example (see also Chapter 6), a blind person who uses a cane for navigation will sense completely different attributes of the environment and relate to completely different landmarks from those encountered by a seeing person [31,58,130].

3.2.3 Precision in Cognitive Maps

How precise are cognitive maps when compared to the environment they describe? To illustrate cognitive maps' precision try to answer the following question: "Which city is located more to the west: Reno, Nevada or Los Angeles, California?" Most people will erroneously answer that Los Angeles is located to the west of Reno [25,58]. We do so since we must classify and cluster the massively detailed cognitive spatial information we face in everyday life. In this case, most of us will hold cognitive maps in which the entire State of Nevada is east of California in a rectilinear manner, which is flawed in reality [25,58].

The need for hierarchy and simplification leads us to gather objects and landmarks into classes and regions in our cognitive maps. The borders of these regions can be geographical but can follow any other objective or subjective criteria that makes sense to us (for example, downtown, west of the river, the poor part of town, the vicinity of my friend's house). Such clusters of landmarks usually follow a hierarchical, multilevel organization [25,58,97]. While we might have geometrically precise survey knowledge of each level, we usually suffer from imprecision when it comes to the geometrical relations between the different levels of our cognitive map hierarchy. For example we might have several geometrically precise cognitive maps, each in isolation, of our house, our neighborhood, downtown streets, our office building and the ring road that we use to go downtown. Attempts to integrate these cognitive maps, each occupying different hierarchal levels, will be highly subjective, inaccurate and prone to error [25,58,97].

Obviously, cognitive maps suffer from imprecision due to temporal changes in the physical environment. Unless updated, cognitive maps are static while the physical world might change in both time and space, which usually leads us to choose stable objects (both physically and temporally) in the environment as our landmarks [58].

Cognitive maps also suffer from geometrical scaling and regularization problems [25,58,68]. Regularization problems are usually manifested in rectilinear, simplified cognitive maps of what is usually a much more complex physical environment (much the same as the Reno-Los Angeles example). Scaling problems appear as consistently compressed or stretched cognitive maps, compared to the physical environment they represent [10,25].

It is important to note that even imperfect and sometimes geometrically erroneous cognitive maps can serve our wayfinding needs. Strictly topological cognitive maps can sometimes serve us flawlessly. For example, most travelers in underground tunnels have a mere topological knowledge of the tunnels layout, and still reach their destination effortlessly [8].

3.2.4 Wayfinding and Virtual Reality

Many researchers have experimented with VR as a wayfinding and navigation training tool (for an overview see [28,31,83]). VR-based trainers promise versatility, compactness, portability and ultimately affordability. They should be particularly useful when the environment to be learned is inaccessible, expensive to explore, dangerous, or imaginary. Reported applications of VR as a wayfinding training tool range from indoor fire fighting simulation to outdoor wilderness and urban navigation simulation [27,28,83]. The army and entertainment industry are especially interested in these applications of the technology.

The VR technology that supports these applications is diverse. We will only mention a few relevant facts concerning this technology. The hardware that is being used for wayfinding trainers ranges from low-end desktops supported by a mouse and keyboard interface to head mounted displays (HMDs), 6 DOF trackers, CAVEs and walking treadmills [27,28,83,156].

The use of tangible mediators as interfaces for better spatial orientation in a virtual environment was also suggested in the Worlds in Miniature (WIM) interface [151]. Much like the Passive Real-World Interface Props (Section 2.6), WIM used a tracked clipboard as a tangible replica of the virtual environment and a tracked ball as a manipulation tool. The user, navigating the virtual environment, could switch to an exocentric 3D-miniature view of the world and manipulate this view by moving or rotating the tangible clipboard and the ball. The WIM can be viewed as a VR elaboration of a 2D map and can improve spatial behavior in the virtual environment [83,151].

The interaction techniques used for navigating the virtual environment and the level of participant control has dramatic impact on the type and quality of the training gained [124]. At one extreme, the participant can be passively moved through the environment; at the other extreme she can be given full freedom and active control of the exploration. Some findings point to the superiority of active exploration over passive exploration (for example [83]), while another shows the opposite trends [131].

Note that the VR technology that is being used for learning a spatial environment does not necessarily have to be visual. Lumbreras and Sánchez [94] created a spatial sound virtual environment they called AudioDoom that can help blind children construct cognitive spatial maps of an imaginary environment (see also Chapter 6).

While it seems obvious that high levels of VR immersion would lead to more effective spatial learning, the work of Patrick et al. [109] has shown that cognitive map trainers based on an expensive stereoscopic HMD and on a less expensive monoscopic large projection screen reached the same levels of effectiveness.

The major concern that overshadows VR-based cognitive map trainer is the problem of knowledge or training transfer.[7] The transfer problem can be demonstrated through a set of questions. Was the cognitive map acquired in the virtual environment useful in the physical world? Was it comparable to cognitive maps acquired in the real world or by other means? Currently there is no clear-cut answer to these questions [27]. While some researchers present encouraging transfer results [83,169], many report mixed or sometimes negative transfer results [27,29,30,57]. One of the most active researchers in this field, Rudolph Darken from the Naval Postgraduate School, presents results from navigation training applications in which VR-based navigation training sometimes actually hindered the development of survey knowledge [30]. Darken claims that while he cannot report that VR was useful as a training aid for navigation, it was very useful as a prediction

[7] The issue of transfer underlies all VR-based training; see for example [124].

tool for the user's navigation ability. Different users have practiced right or wrong strategies consistently in both the VR-based and the physical world, enabling the assessment of their navigation abilities by their performance in the virtual environment [30].

3.2.5 Probing Cognitive Maps

Several techniques are used for the assessment of cognitive maps and their accuracy. All the techniques refer to the physical or virtual world that was perceived by the user as a reference for the accuracy and level of detail of the cognitive map. We will briefly discuss verbal, bearing and distance, map based and functional techniques for cognitive map assessment.

Verbal techniques [9] ask the user to describe the environment through verbal communications. The user is asked to describe a route in an environment or refer to positions of objects in detail. Verbal techniques can achieve insight into the user's cognitive map since the verbal description often employs verbs of motion in addition to the dry report of landmarks' physical locations. Obvious drawbacks of such methods are the subjective and imprecise nature of the verbal information and the need to compensate for different oral abilities. These methods are rarely used as the sole assessment technique for cognitive mapping ability.

The bearing and distance technique is very common in cognitive map assessment [10,15,24,28,57,68,83,133,143,157,168]. The user is placed at a certain location in the environment or is asked to imagine being at a certain location (in which case the user will usually be blindfolded) and is asked to point to another object in the environment and to estimate the distance to it. The distances recorded can be along a theoretical straight line or take into account the physical route traveled between the points [157]. These inter-object distances and directions can be recorded and compared to the distances in the original environment. A spatial map of the environment can be assembled from this data and compared to the spatial layout of the original environment [10]. The distance and direction estimations are egocentric in nature and are usually scaled to a fixed, metric measurement system, although some have used more subjective, non-metric, distance scales [68].

The major advantage of the bearing and distance technique is its ease of implementation—the user does not have to be repeatedly placed in the actual physical environment or to repeatedly navigate her way in it for an assessment to be made [25]. However, the technique suffers from scale problems, since our ability to accurately estimate distances is limited [25,68]. Moreover, the technique will also have low sensitivity to high levels of survey knowledge such as the ability to generate new (not previously traveled) paths [25,27].

Map drawing or map placement techniques ask the user to describe her cognitive map in a spatial manner. In a map drawing assessment the user is requested to sketch a replica of her cognitive map. This technique is sensitive to different sketching abilities and subjective scaling problems and is rarely used [4,25,27,68]. Much more established is the map placement technique. In this technique the user is presented with a fixed grid and asked to point to an object's location or to place representations of an object tangibly on the grid [10,57,68,109,156,157]. Physical objects or small-scale models have also been used to replace map sketching, especially with children

[65,113,114,146]. For example, in "Model Village" Piaget used cardboard models of a church, houses and trees to help children input cognitive maps of a prototype environment [114]. Each cognitive map constructed by the participant was sampled manually in a long process (several-minutes, not including the time required for manual placement error analysis). The tangible-constructional element of such map placement probing techniques was shown to improve the participants' performance and accuracy [65]. It is important to note that both the map drawing and the map placement techniques will suffer from the cognitive map's imprecise characteristics mentioned in Section 3.2.3, namely scale, regularization and clustering.

Last, and probably the most profound, are the functional assessment techniques. These techniques take the user back to the spatial environment she tried to learn and assess her ability to perform a previously unknown navigation task in that environment [15,25,57,169]. A basic task will measure the time it takes the user to walk from a point of entry to a point of exit, counting and measuring the magnitude of errors. In a route reversal task the user will be asked to plan and navigate an unknown opposite route from the exit point to the entry point [57]. The user might also be asked to improvise, or face simulated detours [25]. The functional technique gives an excellent insight into the user's survey knowledge [25,57] but also raises a psychological Heisenberg-like principle: each exposure to the environment for assessment is also another exposure to the spatial layout and another learning opportunity. Hence, the knowledge you are attempting to measure is being altered as you measure it [27].

Although technology is being used for wayfinding training and assessment (see the use of VR in Section 3.2.4), the use of computers for explicit probing of underlying cognitive mapping ability is currently very limited. The first use of a computer for cognitive map assessment dates to the late 1970s when John Baird and his group designed a computerized map placement assessment technique. The technique used a simple computer interface (using arrow and letter keys on a keyboard) for inputting building locations on a 13X13 matrix displayed on a monitor [10]. Later, the bearing and distance technique was automated by enabling the participant to input her estimation for bearing and distance using a mouse and a graphical directional arrow [24,143]. In a few VR-based cognitive map trainers, the user's field-of-view (FOV), hand or pointing device are tracked and used to input bearing assessments directly to the computer [15,133].

3.3 Measuring Constructional Ability and Mental Rotation Ability

3.3.1 Probing Constructional Ability

Muriel Lezak defines constructional functions as *"perceptual activity that has motor response and a spatial component"* [89]. Constructional ability can be assessed by visuoconstructive, spatial tasks that involve assembling, building and drawing. In a typical constructional assessment, the participant is presented with a spatial pattern and is asked to mimic it by manipulating or assembling physical objects [89]. The test administrator scores participant performance using measures such as time-to-completion and accuracy, or more demanding observations such as order of assembly and strategy analysis. As far as we know, none of these tests were ever automated or computerized.

Constructional functions and disorders can be associated with impairments such as lesion of the non-speech, right-hemisphere of the brain and early phases of AD, and can be useful in their assessment [63,89]. Constructional function assessment based on the assembly of physical tangible objects generates assessment tools that are non-verbal, relatively culture-free and can be very sensitive to and selective for constructional ability alone [63]. 2D and 3D constructional tasks have been shown to distinguish between different levels of impairment, suggesting that the more complex 3D construction tasks might be more sensitive to visuoconstructive deficits that were not noticeable on the simpler 2D tasks [89].

2D constructional assessments are widely used. WAIS (see Section 3.1.1) contains two physical construction subtests, Block Design and Object Assembly (Figure 3.2) [63,89]. In the former, the participant arranges red and white blocks to copy a presented pattern. In Object

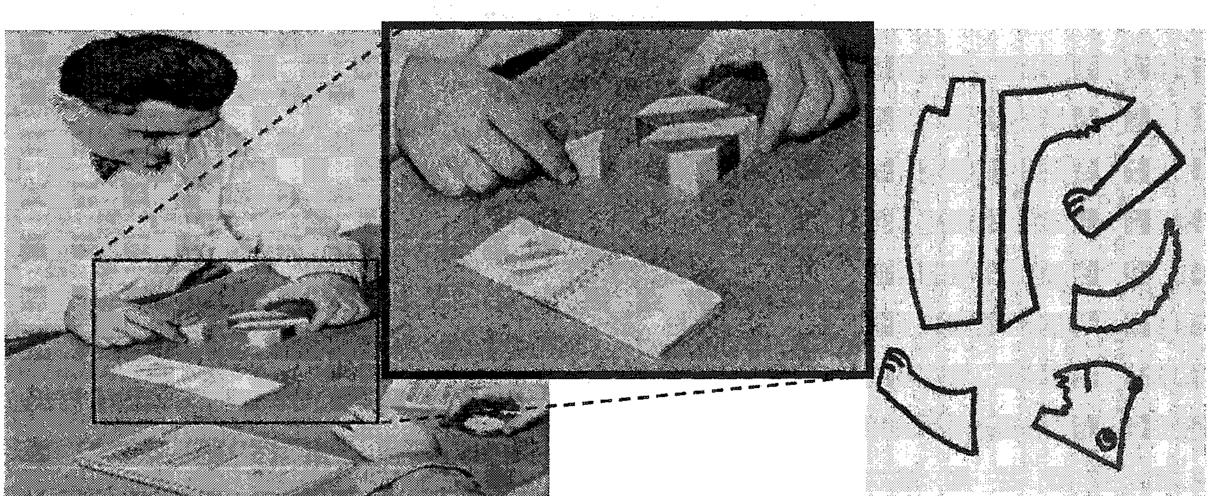

Figure 3.2 – WAIS subtests, from left to right: Block Design; Object Assembly

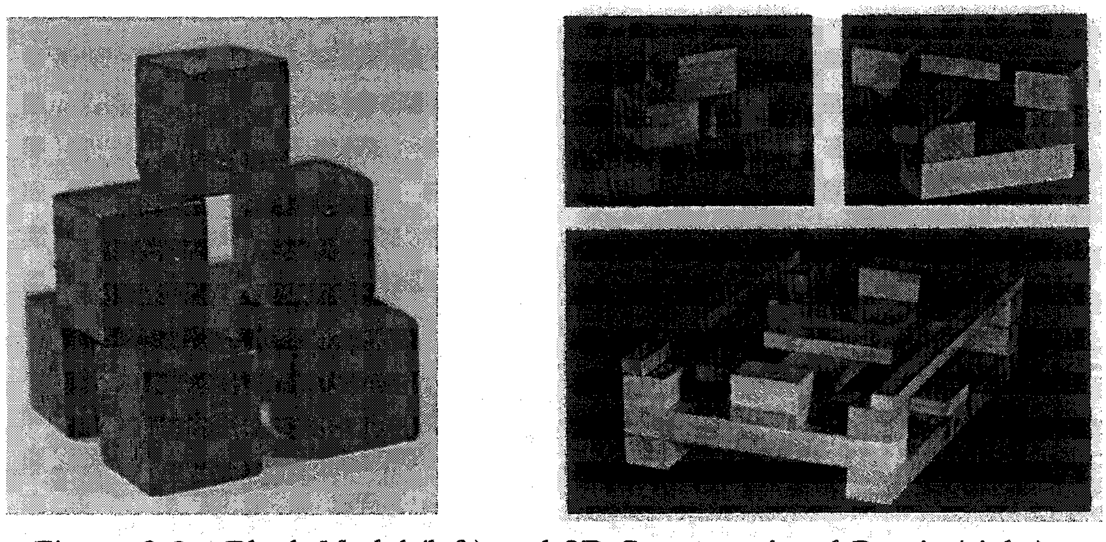

Figure 3.3 – Block Model (left) and 3D Constructional Praxis (right)

Assembly, the participant solves a 2D puzzle. Measures for both tests are based on time and accuracy [63,89].

3D constructional assessments are far less common. Two examples are [89]: Block Model from Hecaen et al. and Three Dimensional Constructional Praxis from Benton et al. (Figure 3.3). In both of these tests the participant tries to match a 3D prototype using wooden blocks, and is scored on time and accuracy. The use of Lego blocks was suggested for 3D tests [89], but to our knowledge was never implemented. Given the complexity of the target shapes in these 3D assessments, manual scoring of even simple measures such as accuracy can be very difficult. Manual scoring of denser measures such as order and strategy would certainly require a very skilled, trained and alert assessor.

3.3.2 Mental Rotation Test

The Mental Rotation Test (MRT) is a paper-pencil based assessment of the visuospatial ability to "turn something over in one's

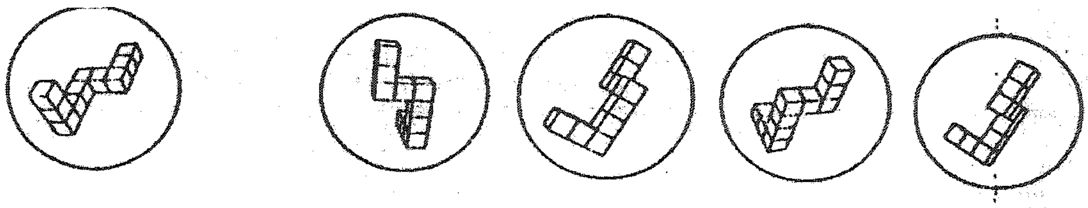

Figure 3.4 – Mental Rotation Test (MRT), sample task

mind" [126,144]. This ability underlies many everyday activities, for example, using a map, or some components of driving [126]. The MRT is 3D and spatial, but in its common form it does not have physical nor constructional components and is purely cognitive.

The MRT is based on early work by Shepard and Metzler [144] that was further established in Vandenberg and Kuse's MRT [167]. MRT's participants are presented with a group of five perspective drawings of 3D objects, one of them is the prototype (the "criterion") object and the rest consists of two identical, but rotated objects, and two "distractor" objects (mirror images of the prototype or simply different objects, see Figure 3.4). The participant is asked to find and mark the two objects that are identical to the prototype object [167].

It was shown that the time needed to determine whether two MRT perspective drawings of objects are similar or not is a linear function of the angular difference between them [144], suggesting that people perform the MRT tasks mentally as if they were physically rotating the objects. The MRT's almost perfect linear relationship between task difficulty and observable human behavior is rare in cognitive assessment; following this relationship the MRT received considerable attention and was extensively researched. The Virtual Reality Spatial Rotation (VRSR) is an automated, VR-based derivative of the MRT [85,126]. In the VRSR participants are asked to manually orient an MRT-like object until it is superimposed on a target prototype. The VRSR adds motoric component and enhances the ecological validity of the MRT by presenting the task in a highly immersive VR environment and by enabling the participants to manipulate the virtual object using a tracked physical prop [85,126].

Chapter IV
Cognitive Cubes

4.1 Introduction

Our research goals were quite simple: to demonstrate that TUIs are practical for resolving meaningful real-life problems, while providing considerable benefits over existing solutions and revealing new possibilities that were not viable without a TUI.

To attain our goals we designed several detailed conceptual mock-ups (see Section 1.3, Chapter 6 and [138]). Each of these conceptual mockups could have potentially become a proof-of-concept, demonstrating the fulfillment of our research goals. Of these mock-ups two were fully implemented and tested: Cognitive Cubes and the Cognitive Map Probe.

In this chapter we will introduce Cognitive Cubes (Figure 4.1), the first half of our proof-of-concept. The second half, the cognitive map probe, will be discussed in Chapter 5. With Cognitive Cubes we designed and tested a specialized TUI for a practical cognitive assessment application. As mentioned earlier (see Sections 1.3 and Chapter 3), our choice of the cognitive assessment domain was a result of balancing our desire to find a practical, useful applied domain for TUIs, our design heuristics (Section 2.2), and the current state-of-the-art of early TUIs. Tackling cognitive assessment challenges allowed us to approach the highly significant real-life problem of constructional ability assessment with current TUI technology.

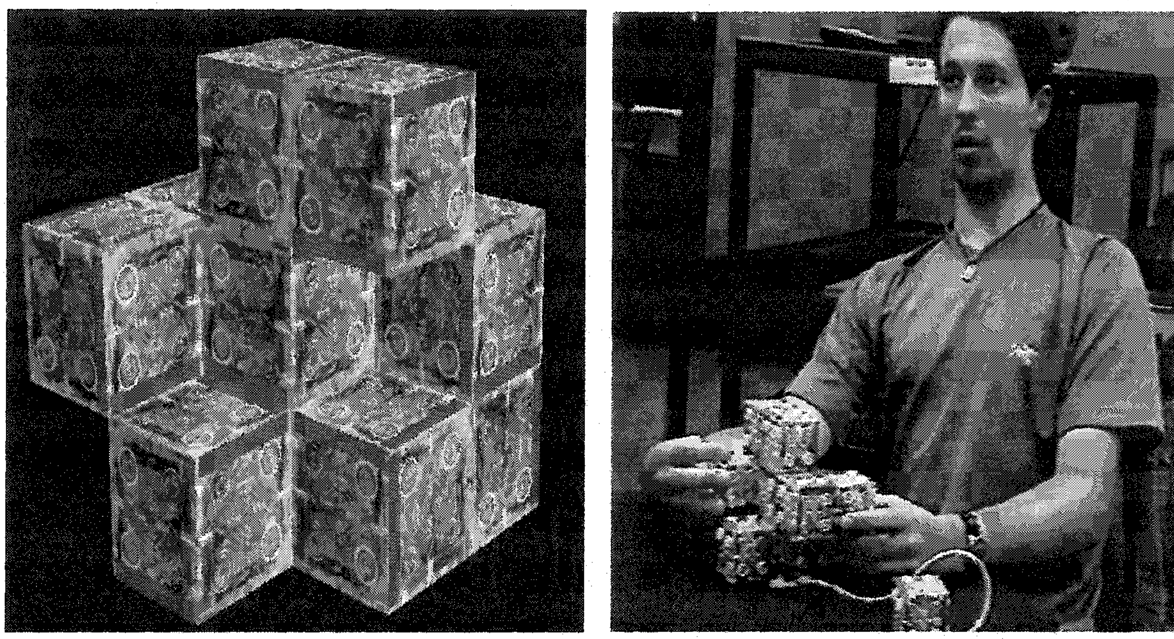

Figure 4.1 – Cognitive Cubes: virtual prototype (left); physical interaction (right)

The Cognitive Cubes theme follows a simple assessment model: show participants a virtual 3D prototype and ask them to reconstruct it physically with a spatial TUI. The prototype presented to the participants is an abstract 3D geometrical shape, constructed of generic-looking building blocks. The TUI consists of a set of identical physical building blocks affording 3D construction, much like Lego blocks.

As we discussed in Section 3.3, the assessment of cognitive spatial and constructional ability is an important clinical tool in the diagnosis and monitoring of brain disease or injury [63,89]. It is also indispensable in scientific study of cognitive brain functions. Techniques for assessment include asking patients or participants to perform purely cognitive tasks such as mental rotation, as well as constructional tasks involving arrangement of blocks and puzzle pieces into a target configuration. These constructional tasks have the advantage of probing not only pure spatial ability, but also the ability to perceive, plan, and act in the world. Studies suggest that assessment with 3D forms of these tasks may be most demanding and sensitive [89]. However, use of 3D tasks in assessment has been limited by their inherent complexity, which requires considerable examiner training, effort and time if scoring is to be consistent and reliable.

Cognitive Cubes was designed as an automated tool for examination of 3D spatial constructional ability. Cognitive Cubes makes use of ActiveCube [78-80], a Lego-like tangible user interface for description of 3D shape. With Cognitive Cubes, users attempt to construct a target 3D shape, while each change of shape they make is automatically recorded and scored for assessment.

We created Cognitive Cubes closely following our TUI design heuristics (See Section 2.2). First and foremost, Cognitive Cubes offers a very intuitive spatial mapping (Section 2.2.2) between the TUI and the assessment task. Most of the constructional assessment activity is performed entirely in the physical domain, using the physical cube-based TUI which, much like Lego blocks, naturally affords constructional activity. The assessment task involves the presentation of a virtual 3D prototype that the participant attempts to physically reconstruct. We kept the virtual prototype in close visual agreement with the physical cubes, texturing it with a detailed matching texture, sampled from the physical cubes (see Figures 4.1-4.4).

At first glance, Cognitive Cubes do not offer strong I/O unification (Section 2.2.3) because the virtual prototype is presented separately from the physical interface. This argument would have been true if Cognitive Cubes were used for 3D design, however, Cognitive Cubes was used for cognitive assessment. The prototype presented to the participant is merely a visual representation of the cognitive goal the participant is expected to reach, and in this sense the prototype is external to the interaction. A tighter coupling between the presented prototype and the physical TUI would leave very little challenge in the constructional task. We argue that Cognitive Cubes offer good I/O unification since the input—actions performed on the 3D cubes, fully coincide with the output—virtual 3D shapes registered at the host computer.

Lastly, Cognitive Cubes, like many other construction sets, offers extremely flexible exploration of the design domain and trial-and-error actions (Section 2.2.4). Participants can perform actions on the 3D structure in any desired order, undoing their former actions in a completely nonlinear fashion (that is, undoing actions in an order that does not follow the construction order).

As far as we know, Cognitive Cubes is the first computerized tool for constructional assessment, combining the increased sensitivity of 3D constructional tasks with the efficiency, consistency, flexibility and detailed data collection of automation. In this chapter we detail Cognitive Cubes hardware and software and present a methodical experimental confirmation of the sensitivity and utility of our applied TUI.

4.2 Hardware

4.2.1 Infrastructure – ActiveCube

To measure 3D constructional abilities we needed an interface that will maintain a high level of 3D physical constructional expressiveness while enabling precise real-time sensing of the structure's geometry. Our early aborted attempts to design our own 3D TUI, BLOXELS, in 1999 were mentioned in [138] Later in 1999 we tried unsuccessfully to borrow the Blocks TUI from MERL for the same purpose (see Section 2.11, Figure 2.38).

ActiveCube (see Section 2.11, Figure 2.40) is probably the best current example of a spatial 3D TUI for structural input. ActiveCube was developed by Dr. Yoshifumi Kitamura and his group, part of Fumio Kishino's lab at Osaka University, Japan. One of the dominant developers of the ActiveCube software and hardware was Yuichi Itoh, at the time Dr. Kitamura's Ph.D. student. When we designed Cognitive Cubes, ActiveCube was the only interface that enabled real-time, step-by-step, geometry sampling of a 3D structure (for comparison, sampling of a single, elaborate structure built with MERL's Blocks 3D TUI can take almost an hour, see Section 2.11). We were fortunate to meet and initiate an extremely fruitful collaboration with Dr. Kitamura and his group in early 2000. Cognitive Cubes are the fruits of this collaboration. The adaptation of ActiveCube to our Cognitive Cubes design was mostly performed by Yuichi Itoh, who later came to Edmonton to assist us during the two months of the Cognitive Cubes experiments.

ActiveCube consists of a set of plastic cubes (5 cm/edge) that can be attached to one another using male-female connectors (employing simple clothing-like snaps), forming both a physical shape and a network topology. Each cube and cube face has a unique ID. A host PC is connected to a special base cube and communicates with the small CPUs in each cube through a broadcast mechanism to sense the (dis)connection of any cube. Since all cubes have the same size and shape, any topology represents a unique collective shape [78-80].

As we mentioned in Section 2.11, ActiveCube capabilities include more than 3D geometry input. The cubes are equipped with a large variety of input and output devices, supporting flexible interaction paradigms. Some of the cubes are equipped with ultrasonic, optical (visible and IR), tactile, gyroscopic and temperature sensors. Other cubes are equipped with light, audio, motor and vibration actuators [70,78-80].

4.2.2 Customizing Cognitive Cubes

To support our constructional ability assessment paradigm, and to allow us to assess participants with diverse constructional abilities (we were planning to approach young, elderly, and participants with mild AD), Cognitive Cubes hardware had to support the following functions:

1. Allow flexible 3D geometry input by assembly of physical cubes.
2. Sample the physical 3D cubes structure in real-time.
3. Allow easy handling of the hardware. Cubes had to be connected to, and disconnected from, each other in a straightforward manner.

To accomplish these requirements, Cognitive Cubes needed only a subset of ActiveCube capabilities, namely the interactive 3D geometry inputting. In this sense, ActiveCube additional input and output capabilities could well be distracting for Cognitive Cubes purposes. We decided to work only with a generic ActiveCube, using cubes with the same color and shape, without any of the extra ActiveCube sensors or actuators.

To ease the connectivity of the cubes we added a blue stripe on each of the cubes faces. To snap the connectors for proper assembly required that the user either match the male-female connectors, or match the two blue stripes on the two connecting faces (see Figures 4.1 and 4.2 for Cognitive Cubes appearance).

4.3 Software

4.3.1 Facilitating Prototype Viewing and Interaction

To assess constructional ability using Cognitive Cubes, we needed to present the participant with the prototype she is asked to construct. There are several possible ways in which a participant can view and learn the prototype. Let us briefly discuss these options and the choices we made:

1. Virtual vs. physical: The prototype can be a physical object or a virtual entity, presented to the participant using a display method. We believe that physical prototypes presentation can offer an interesting approach to constructional ability assessment [134]. The participant can freely view, touch, and rotate such a prototype, while trying to reconstruct it using Cognitive Cubes. The physical prototype can be constructed in different scales than the TUI cubes, and can appear as a single solid object, rather than an assembly of cubes. It seems that a physical approach to prototype presentation might afford easier shape learning, and perhaps allow participants to successfully engage more complex tasks. Working on a first-of-its-kind device we decided against the physical approach because of the extra resources it required and the limited flexibility it offered (physical prototypes have to be constructed and once they are built editing them can be cumbersome). We should

also mention that prototypes can be presented as shapes drawn on paper. This option raises the problem of presenting all aspects of the 3D prototype on the 2D paper. While several solutions come to mind (for example, drawing hidden lines, drawing different aspects side by side, working only with prototypes that expose themselves through a single aspect, etc.) they all use high levels of abstraction. We believed that by introducing such abstraction we might "lose" participants, especially among the elderly or participants with mild AD, who might fail to visualize the prototype correctly before attempting to build it. We chose a virtual display as our prototype presentation method. While virtual displays impose a certain level of abstraction (the virtual object is not really there), they can offer relatively high levels of realism and afford an extremely flexible prototype presentation, enabling us to test and edit easily the vocabulary of our prototype shapes.

2. The means of virtual display can vary dramatically from stereoscopic VR HMDs and CAVEs displays to monoscopic projectors and screens. As discussed earlier (Section 3.2.4, [109]), higher immersion doesn't necessarily mean better interaction. Adding immersion through an HMD or CAVE shutter glasses might hinder prototype learning, especially with our elderly and AD participants who might find the VR interface obtrusive, preventing us from measuring valuable data in the TUI-based assessment phase [90]. We chose to project the 3D virtual prototype in front of the participant (see Figure 4.2) using a monoscopic digital rear projector and a large screen (125cm diagonal).

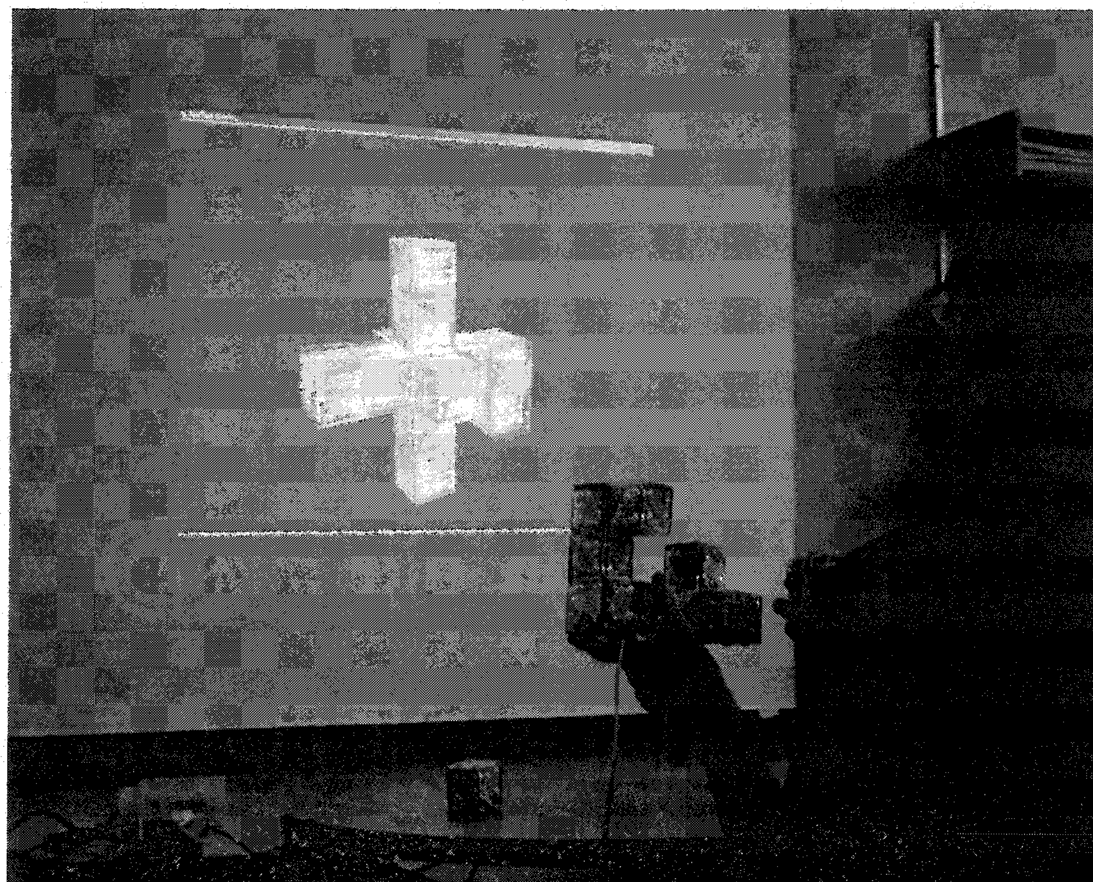

Figure 4.2 – Cognitive Cubes: prototype display and interface

3. After selecting a virtual prototype display we had to support the participants viewing all of the prototype aspects. An interesting viewing scheme could employ an attached 3 DOF orientation tracker to the physical base cube, enslaving the virtual model to the tracker. In this manner the participant can choose the displayed prototype aspect simply by orienting the physical structure. Another, simpler, option is to continuously rotate the virtual prototype at a constant pace around its vertical axis, providing 3D depth information. The rotating prototype option also engages the participant in mental rotation (see Section 3.3.2) and use of memory as the virtual prototype and the physical cubes

orientations match only periodically. Seeing the benefits in both approaches, we chose the rotating virtual prototype, weighing its implementation simplicity and the enhanced mental rotation flavor it added to the Cognitive Cubes assessment tasks. After several iterations (see Section 4.6) we fixed the prototype rotation at a slow 2.7 rpm (revolutions per minute) pace.

4. Working with a virtual prototype display we could also consider enhancing the memory component of our task by using a vanishing prototype. In this manner, the participant would have a chance to learn the rotating prototype for a limited time. After that period passes the prototype would vanish from the display and the participant would have to reconstruct the shape using only their memory. We decided to experiment first with a non-vanishing presentation and found out (Section 4.6) that the remaining constructional tasks are sufficiently difficult, and challenging for participants.
5. To add realism to the virtual prototype, each virtual cube face was textured with an image of a physical cognitive cube face (see Figures 4.1-4.4).

Other than the display, Cognitive Cubes software supported interaction with minor audio cues: when the participant connects a cube to the structure a distinct chime sounds through a speaker. If the participant chooses to disconnect a cube, a different chime sounds.

During an experiment, the assessor could easily switch between virtual prototypes using a simple menu tool. The software did not generate any cues about the precision of the physical Cognitive Cubes structure vis-à-vis the virtual prototype. Hence the participant worked freely and when satisfied with the match between her construction and the prototype, she informed the assessor, who advanced the system to the next trial. The assessor could also choose to stop the assessment at any point if, for example, the participant was not making any progress.

4.3.2 Probing

While the participant attempts to reconstruct the virtual prototype Cognitive Cubes collects a data vector, containing the following values, for each participant action:

1. Event time: in seconds, measured from the time the virtual prototype appeared on the display.
2. Action type: cube connection or disconnection.
3. Cube location: can be viewed as a Cartesian set of coordinates, measured from the base-cube which is located at the origin.

After assessment the collected data is analyzed offline to calculate the 3D similarity between the participant's structure s and the prototype p. *Similarity* is calculated for each connect or disconnect event. For example, a five steps participant assembly will result in five different *similarity* calculations. The equation for *similarity* is presented in Equation (4.1), where i is an intersection of s and p, and $|i|$, $|s|$, and $|p|$ are the number of cubes in i, s and p. s is maximized over all possible intersections i produced by rotating or translating s. Intuitively

$$(4.1) \qquad Sim = 100 \cdot \left(\frac{|i|}{|p|} - \frac{|s| - |i|}{|p|} \right)$$

speaking, *similarity* is the number of intersecting cubes minus the number of remaining "extra" cubes in the participant's structure, normalized by the number of cubes in the prototype.

We make *similarity* at task completion, calculated as described above, one of our four assessment measures. The remaining three are: *last connect*, the time elapsed from the start of the task to the last cube connect or disconnect; *derivative*, the differences between two successively measured similarities in a task divided by the time elapsed between those measurements (local "slope" of the similarity function), averaged for all such pairs in a task; and *zero crossings*, the number of times the local slope crossed zero. We sometimes use the terms "completion time", "rate of progress", and "steadiness of progress" as substitutes for *last connect*, *derivative*, and *zero crossings* [137].

4.4 System Strength and Weaknesses

In view of the assessment task it was designed for, Cognitive Cubes suffered from a few technical limitations. It is important to stress that as a pioneering prototype, Cognitive Cubes limitations should be viewed only in comparison to a hypothetical ideal TUI-based assessment device.

Cognitive Cubes is essentially a construction set. Thus the cubes will constantly be connected and disconnected. The simple cube connectors not only attach the cubes to each other, they also have electrical functions (distributing power and information, see Section 2.11). We found that the wear-and-tear on these connectors is quite considerable and that mechanical connectivity problems soon translate into discrepancies in the electrical behavior of the cubes (see Section

4.6.4 for a discussion of the different Cognitive Cubes errors and their rates).

With Cognitive Cubes, any cube face can be attached to any other cube face, unlike Lego blocks, which can only be stacked one on top of another. This design freedom comes with a price and can lead to unstable arrangements since the cubes' connectors can only support the weight of a few cubes. For example, an unsupported "arm" of several cubes reaching out horizontally is prone to collapse. Although we could ensure that prototypes were well-balanced shapes, we could not prevent participants from building unstable structures.

Last, a few participants found aligning the cubes for connection (see Section 4.2.2) difficult and required a longer training period before being able to approach the assessment tasks. Some participants also used unnecessary force when connecting the cubes, perhaps due to lack of confidence in their alignment of the connectors. This extra force manifested itself in connector deterioration and, ultimately, errors.

Cognitive Cubes offers many dramatic advantages over existing tools for constructional assessment. These include:

Consistency. Any assessment is of little use if comparisons between different sets of its results are not trustworthy. Existing tools for 2D assessment have been in use for some time and are quite consistent (See Section 3.3.1). However, while 3D constructional assessments have been proposed previously (see Section 3.3.1), the complexity of the shapes and tools involved have made them difficult to administer consistently. As the first automated constructional assessment, Cognitive Cubes is extremely consistent in administration and scoring.

Sensitivity. Assessments must also respond as sensitively as possible to the strength or weakness of cognitive abilities. Because they incorporate a demanding level of complexity, 3D constructional assessments have shown particular promise in this regard. Cognitive Cubes incorporates this sensitivity. Moreover, automation allows Cognitive Cubes to monitor elements of performance that are ignored in other assessments, including actions during the assessment itself (not just the final result).

Cost efficiency. Although the hardware used by Cognitive Cubes will be more costly than that used in many other assessment tools, we expect that overall it will reduce costs. Because Cognitive Cubes is automated, the level of training and expertise required by personnel employing it will be relatively inexpensive. In addition, automation should ultimately allow adaptive testing, identifying the level of cognitive ability much more quickly and thus reducing testing time.

4.5 Experimental Methodology

4.5.1 Overview

Is Cognitive Cubes as sensitive and consistent as we would like to believe? In the following sections we describe our experiments, designed to find answers to this question. In this section we begin with a discussion of our general experimental approach. Section 4.6 describes the testing and resulting adjustment of this approach in a pilot study. Next, in Section 4.7 we describe the cognitive sensitivity study, which examined the response of Cognitive Cubes to known participant and task cognitive factors. Finally, in Section 4.8, we present the test comparison study, which compares Cognitive Cubes' results to those obtained with a standard paper-and-pencil 3D assessment. The experiment was conducted with the approval of the University of Alberta Health Research Ethics Board (HREB-B, file number B-020501-COMSCI). See Appendix A for the experiment information letter, consent form and detailed protocol.

4.5.2 General Methodology and Design

Cognitive Cubes tasks were designed with two principles in mind. First, tasks should be as diverse and as interesting as possible, ensuring that participant interest remains high and that the assessment is sensitive to a range of cognitive ability levels. Second, participants should move gradually from easy to difficult tasks. This eases participants into a familiarity with the interface, and allows quick identification of the participant's cognitive ability threshold, permitting participants to drop out without frustration as soon as they show the limit of their capabilities.

We designed four task types (see Figure 4.3). *Intro* tasks were simple practice trials, designed to introduce the participant to Cognitive Cubes. A cube appears on the display after each new connection, indicating the next cube to attach. The *follow* task type also provided step-by-step guidance, but the tasks were much more difficult. *Match* tasks provided no cube-by-cube guidance, but rather displayed a complete virtual prototype for the participant to construct using their own approach. In all three of these task types, the starting point for the participant's construction was the base cube. With *reshape* tasks the participant started from a more complex initial condition (always the same 7-cube 3D construct, see Figure 4.3)—in all other respects reshape was exactly like match.

Since *intro* and *follow* guided the participants, these tasks required less cognitive planning than *match* and *reshape*. We expected *reshape* tasks to require more planning effort than *match* tasks since it started from a complex, somewhat arbitrary shape. For this reason task types

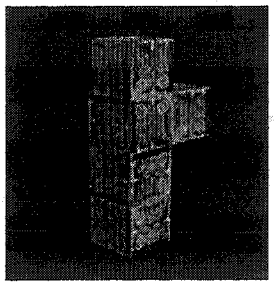

Intro task

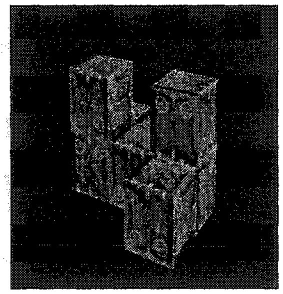

Follow task

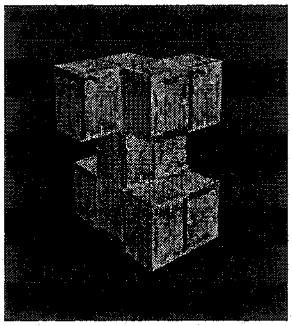

Match Task

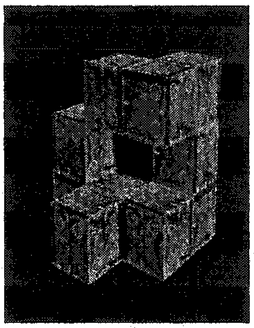

Reshape Task

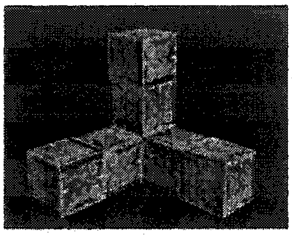

Reshape – initial condition shape

Figure 4.3 – Cognitive Cubes task types (samples)

were placed in an *intro, follow, match* and *reshape* order. Within each task type, tasks were organized roughly according to their relative difficulty, taking into account the number of cubes, symmetry, and 2D or 3D shape. We assessed the complexity of a given virtual prototype by the number of cubes in its structure and whether it was 3D or 2D. By 2D we mean structures that are planar, with all the cubes existing on a single surface (see for example the *intro* task in Figure 4.3). By 3D we mean structures that are not planar (see for example the *follow, match* and *reshape* tasks in Figure 4.3).

Intro shape complexity was minimal, while *follow* tasks used moderately complex shapes. In *match* tasks, shapes reached their greatest complexity, while shape complexity was moderated for reshape in light of the heightened demands on cognitive planning.

During each Cognitive Cubes assessment, a participant performed 39 tasks. The tasks were performed in this order: 6 *intro* tasks, 8 *follow* tasks, 15 *match* tasks and 10 *reshape* tasks. Overall, 14 tasks were 2D and 25 were 3D. The number of each kind of task was picked in order to balance the different assessment phases, task types, level of difficulty and overall assessment time.

4.5.3 Setup and Procedure

The participant sat at a table with only Cognitive Cubes placed in front of her. A 125 cm diagonal image was displayed in front of the viewer at a viewing distance of 185 cm using a digital projector, in a brightly lit room (Figures 4.1 and 4.2). The assessment administrator sat at an adjacent table with the host PC (see Appendix A).

The experiment was conducted with a strict written protocol read out loud to the participant. The participant was introduced to the system, the experiment, and its purpose, and then read an information letter. She was told that she might stop the experiment at any time, and asked to sign a consent form. The participant was given a short interview, answering questions concerning age, education, occupation, experience in 3D design, construction sets, computer games, general health and handedness.

The participant was asked to be "as fast and as precise as possible", and was told that "the system is recording your actions". She was told that there was no time limit, and that she may decide when she had finished each task, but that she should do "the best she could" in building each shape. She was asked to connect one cube at a time to the cube structure (avoiding offline interaction), unless reconnecting a chunk of the structure that had fallen off. In contrast, removing several cubes at once was perfectly fine.

The system never provided feedback about construction correctness. In the first few *intro* tasks the participant was guided closely by the administrator, both verbally and physically. Guidance was gradually reduced and after the *intro* tasks, ceased. Between the *follow* and *match* tasks the participant was asked to take a short break. During the *reshape* tasks, planning was encouraged by reminding the participant that "the system is counting your steps". Finally, the participant was interviewed for her impressions of the system and the experiment. Performing the complete assessment took roughly 90 minutes on average (see appendix A for Cognitive Cubes experiments protocol).

If the system hardware failed, the participant was asked to repeat the interrupted trial. When the participant was not showing progress in a task after 5 minutes, the administrator suggested ending the task. Skipped tasks tended to be the more difficult ones.

4.6 Pilot Study

4.6.1 General

We performed an unexpectedly extensive pilot study due to a few surprises as well as typical need to gain some practical experience with the system and tune it accordingly. The main lessons learned from the pilot study included understanding and tuning task difficulty levels, understanding the mechanical-physical attributes of the virtual prototypes we asked the participants to construct, the need to deal with several unexpected system errors, and the finalization of the experimental protocol.

4.6.2 Participants

Our pilot study included 14 young, healthy participants who performed the entire cognitive assessment. They ranged in age from 22 to 43, with average age of 29.21 and standard deviation of 5.45 years. Of the pilot study participants 3 were females and 11 were males. Most were recruited within the University of Alberta Computing Science Department students and their families and other acquaintances.

Most experimental adjustments were made in response to feedback from the first two participants. Tasks for the remaining pilot study

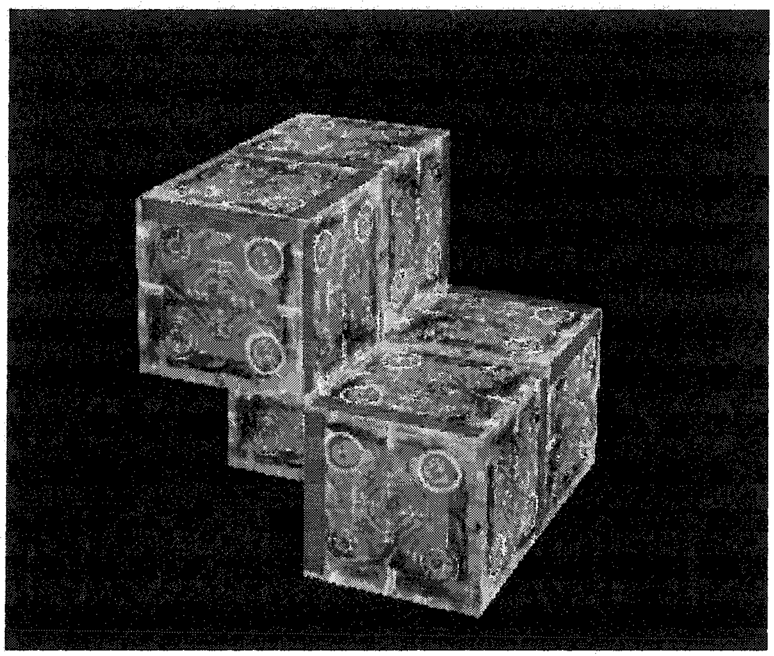

Figure 4.4 – A five-cubes *follow* task

participants were not changed, and their feedback resulted only in some fine tuning of the written protocol.

4.6.3 Results—General

Perhaps our most important lesson from the pilot study, related to the difficulty of Cognitive Cubes. We found that many of our healthy, young participants faced difficulties with tasks that involved a relatively small number of cubes in a 3D arrangement. For example, several of the pilot study participants found the seemingly simple, five-cubes *follow* task in Figure 4.4 quite challenging, though eventually manageable. Matching a ten-cube prototype proved to be very challenging for several participants. Consequently, we decided that all shapes would be restricted to at most ten cubes.

We quickly found that the cubes' ability to support their own weight was limited. Certain prototypes were modified in an attempt to prevent potential collapse of unsupported cubes. Our decision to limit the structure size to a maximum of ten cubes also helped in creating more structurally sound prototypes.

A strict written protocol also emerged from the pilot study. A few task parameters, most notably the prototype rotation rate, were altered early in the pilot study following participant feedback.

4.6.4 Results – Measured Reliability

During the pilot study we learned that Cognitive Cubes suffers infrequently from three types of system errors. The most severe was the connection error, where the system reported cube (dis)connections that did not in fact occur. These events were always excluded from our analyses. The less severe "crash errors" occurred when the system simply stopped responding. We decided to allow the participant to repeat tasks with crash errors. Finally, the cube construction sometimes collapsed, usually when the participant applied too much force. The participant typically reconnected the collapsed cubes immediately without administrator intervention. We filtered for the collapse errors by locating multiple cubes simultaneously disconnecting and within 10 seconds, simultaneously reconnecting. As a result, these errors did not affect the similarity function.

Table 4.1 summarizes the overall Cognitive Cubes errors, as measured in our subsequent experimental phase, the cognitive sensitivity study (see Section 4.7). Note that the cognitive sensitivity study also included elderly and AD participants, who used the system in a fashion that generated more system errors.

Table 4.1 – Cognitive Cubes error trials, occurrence and frequency

Trial Description	**Frequency**
Total trials:	621 (100%)
Error trials:	
Connect errors	26 (4.2%)
Sys crashes w/o connect error	13 (2.1%)
Trials without error:	582 (93.7%)
Remaining affected trials:	
Repetitions	21 (3.6%)
Filtered collapses	81 (13.9%)

4.7 Cognitive Sensitivity Study

4.7.1 General

To confirm and improve the sensitivity of Cognitive Cubes, we studied its response to three factors known to correspond to differences in cognitive ability: participant *age* (≤34, ≥54), *task type* (*follow*, *match* and *reshape*), and *shape type* (2D, 3D).

Since cognitive ability declines gradually with increasing age, in this study we expected younger participants to perform better than older participants. As the cognitive load of a task increased, cognitive abilities are stressed, leading us to expect better performance with *task types* that required less planning. Similarly, we have already noted the heavier cognitive demands involved in working with 3D shapes. We anticipated better performance with 2D shapes than with 3D shapes.

Because cognitive ability decreases with the progression of AD, we made a preliminary study of the sensitivity of Cognitive Cubes to that form of dementia. As we mentioned in Sections 3.1.4 and 3.3.1, there is no existing cure to AD, but its early detection can have an enormous impact on palliative care and quality of life [91]. Although the numbers of AD patients in this study were small, we include some limited results here.

4.7.2 Participants and Methods

The cognitive sensitivity study included 16 participants, recruited on and off campus, ranging in age from 24 to 86, with 4 females and 12 males. 7 of the participants were young, 7 elderly and 2 were elderly with mild AD. The average age of the young was 26.71 with standard deviation of 4.92 years. The average age of the elderly was 70.71 with standard deviation of 10.17 years. The two mild AD participants were 58 years old and 77 years old, both males. In addition to the methods described in Section 4.5 participants also performed a Mini-Mental State Evaluation [49].

4.7.3 Single Task Example

Figure 4.5 provides an informal view of some of the study results. The figure presents the *similarity* (Equation 4.1) versus time, for a single Cognitive Cubes task. The task is a seven-cube, 3D *match* task. The *similarity* measure curves are plotted for the 13 cognitive sensitivity study participants who performed the task.

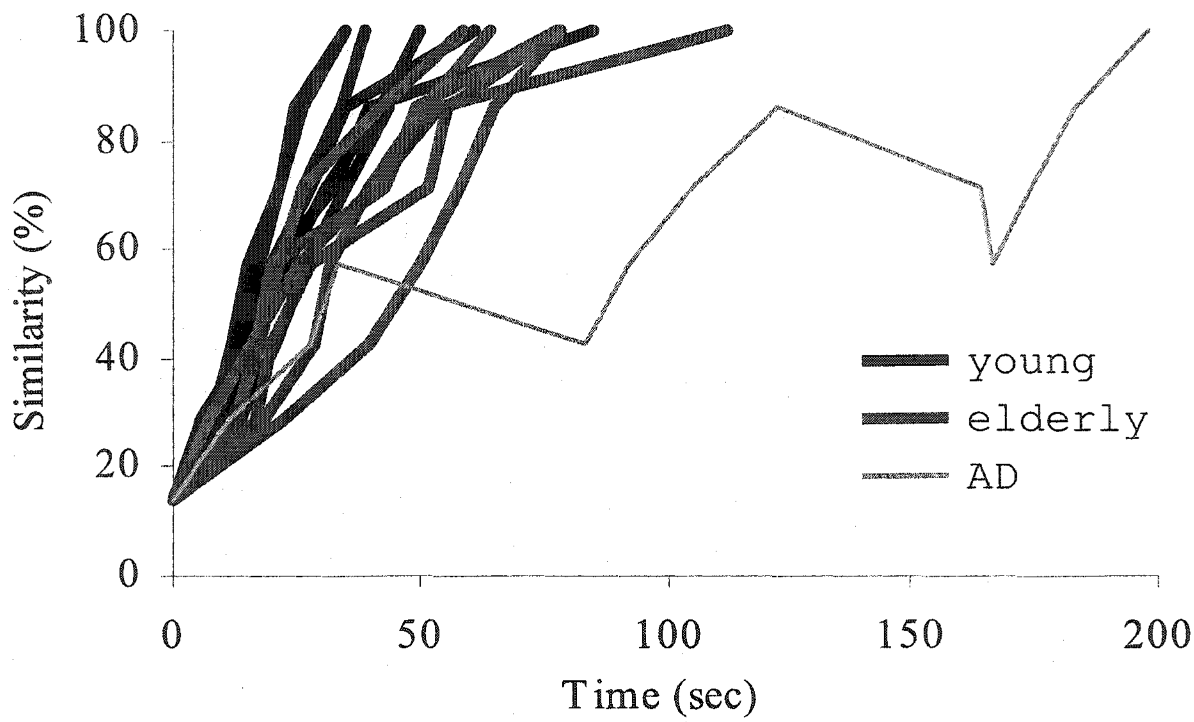

Figure 4.5 – Similarity vs. time; cognitive sensitivity study; single task

It is interesting to note that all the participants who began this task completed it, reaching a final *similarity* of 100%. The *total time* measure (see Section 4.3.2) varies considerably between the different groups: most of the young participants completed the task faster than most of the elderly participants. All of the participants accomplished the task more quickly than the single AD participant.

Participants' rate of progress, as manifested through the curve slope, or the *derivative* measure, also differs between the groups. Most of the young participants have a steeper similarity slope than the elderly participants. All the participants have a faster rate of progress than the AD participant.

Lastly, with the exception of the AD participant, all participants have steady progress towards the goal, and thus no *zero crossings*

(Section 4.3.2). However, the AD participant similarity curve local slope crosses zero several times.

4.7.4 Results

We start by mentioning a few randomly selected qualitative results, based on the post-session interview with the participants:

- Female, healthy, 20—“Fun! Feels likes playing, not a test”.
- Male, healthy, 28—“Rotation is hard, should be controllable”.
- Male, healthy, 73—“Enjoyable, but tiring, too many tasks”.
- Male, mild AD, 58—“Hard to focus, connecting cubes is difficult”.

In the results that follow, we exclude connection errors and system crashes, and filter collapse errors as described in Section 4.6.4. During the experiments, we repeated trials with system crashes and included them in the analyses. Error and repetitions frequencies are listed in Table 4.1.

Because there were so few AD participants, we exclude them from any analyses of variance. Elderly participants often were not able to complete all tasks, unbalancing the age factor in our ANOVAs (as mentioned in Section 4.5.2, *task type* and *shape type* are already unbalanced in the design). The number of trials completed is listed in Table 4.2 (the table also details the potential number of trials, in an ideal case where no trials were skipped or withdrawn).

Results are presented in Tables 4.3 and 4.4 and Figures 4.6-4.9. We analyze the results using one 3-factor (2 *age* x 3 *task type* x 2 *shape type*) unbalanced ANOVA for each of the *last connect, similarity, zero crossings*, and *derivative measures*. We exclude the *intro task type* and the AD participant results from the analyses.

All three factors produced main effects in line with our expectations. Participant performance varies significantly by age (Table 4.4), with elderly participants needing more time to complete each task, and showing a low rate of progress. By all four dependent measures, participant performance is also significantly affected by *shape type* (Figures 4.6-4.9). 2D shape construction is completed more quickly, more accurately, and with a higher and steadier rate of progress. Finally, *task type* also has significant effects on all four measures. *Follow* is the easiest of the task types, enabling quick completion and a high, steady rate of progress toward the target shape. However, shape *similarity* is lowest with the *follow* task. Participants perform the *match* and *reshape* tasks with roughly equal completion times and similarities, but the rate of progress in the *match* task type is higher and steadier.

The only significant interaction, by all four measures, is between *shape type* and *task type* (Figures 4.6-4.9). In general, for 2D tasks, *follow* and *match* are roughly equal in difficulty by all four measures. *Reshape* is more difficult. For 3D tasks, *follow* is simplest, followed by *match,* then *reshape. Last connect* time is the lone exception among the measures: *reshape* tasks are completed more quickly than *match* tasks.

Table 4.2 – Completed trials in the cognitive sensitivity study by participant group, *shape* and *task type* (potential number of trials in brackets)

Group	Overall	Shape Type		Task Type			
		2D	3D	*intro*	*follow*	*match*	*reshape*
Young	270 (273)	97 (98)	173 (175)	41 (42)	56 (56)	103 (105)	70 (70)
Elderly	246 (273)	97 (98)	149 (175)	42 (42)	56 (56)	88 (105)	60 (70)
AD	63 (78)	27 (28)	36 (50)	12 (12)	15 (16)	24 (30)	12 (20)

Table 4.3 – Results of three way ANOVA in cognitive sensitivity study. Intro trials and AD participants are excluded.

Independent Measures	Dependent Measures	ANOVA
age	*last connect*	$F(1,13)=23.82$, $p<.00005$
age	*derivative*	$F(1,13)=71.21$, $p<.00001$
task type	*last connect*	$F(2,26)=4.7$, $p<.01$
task type	*similarity*	$F(2,26)=4.96$, $p<.01$
task type	*zero crossings*	$F(2,26)=7.58$, $p<.001$
task type	*derivative*	$F(2,26)=34.32$, $p<.00001$
shape type	*last connect*	$F(1,13)=37.24$, $p<.00001$
shape type	*similarity*	$F(1,13)=3.9$, $p<.05$
shape type	*zero crossings*	$F(1,13)=13.07$, $p<.0005$
shape type	*derivative*	$F(1,13)=137.15$, $p<.00001$
ttype x stype	*last connect*	$F(2,26)=3.22$, $p<.05$
ttype x stype	*similarity*	$F(2,26)=6.02$, $p<.005$
ttype x stype	*zero crossings*	$F(2,26)=3.93$, $p<.05$
ttype x stype	*derivative*	$F(2,26)=4.16$, $p<.05$

Table 4.4 – The main effect of *age* on *last connect* and *derivative*; Table presents means, with standard error in parentheses

Dependent Measures	**Age Group**		
	young	**elderly**	**AD**
Last connect (sec)	48.45 *(3.5)*	76.33 *(5.3)*	91.24 *(8.8)*
Derivative (sim/sec)	3.02 *(.13)*	1.97 *(.1)*	1.43 *(.16)*

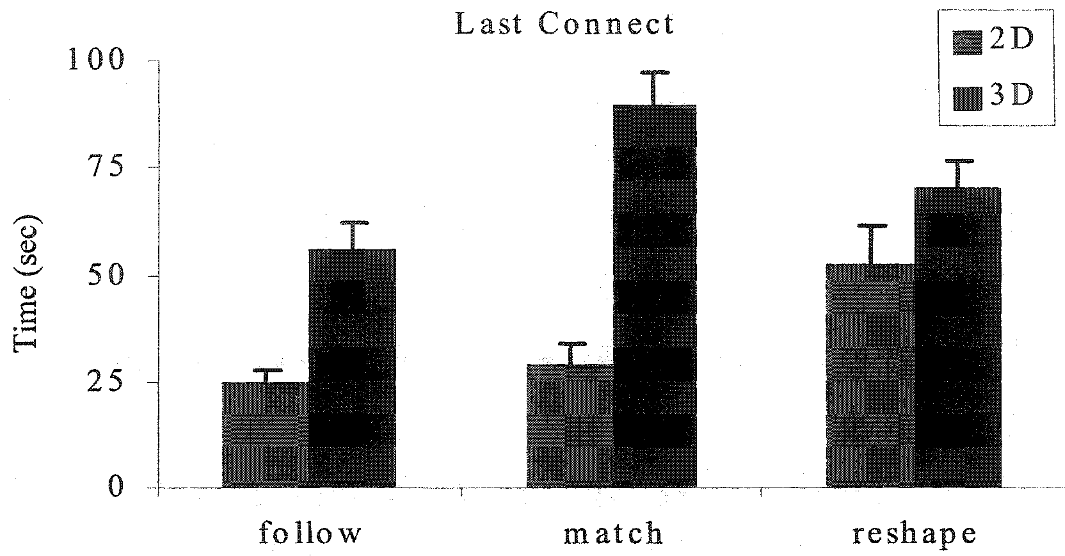

Figure 4.6 - The effect of the *shape x task type* interaction on *last connect*

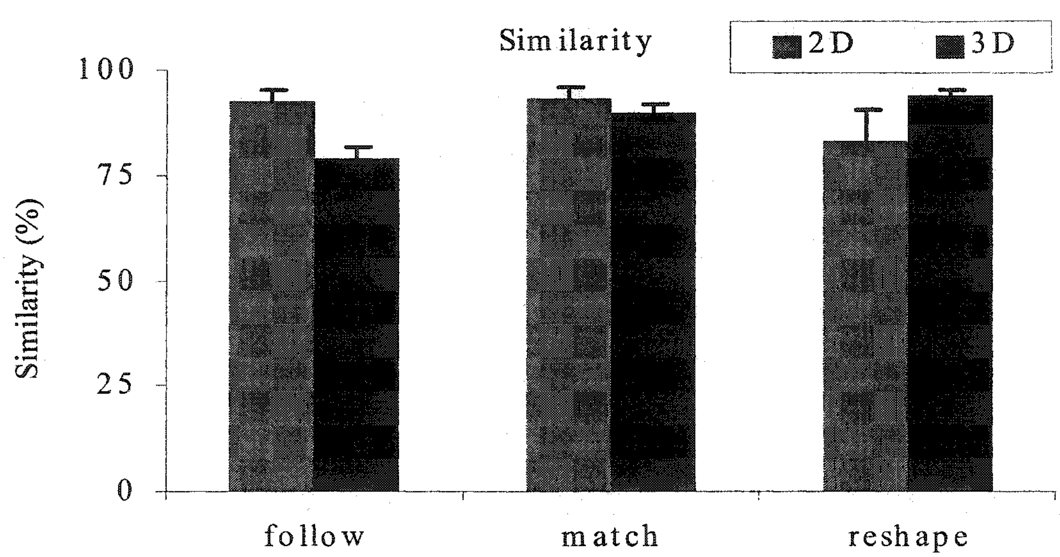

Figure 4.7- The effect of the *shape x task type* interaction on *similarity*

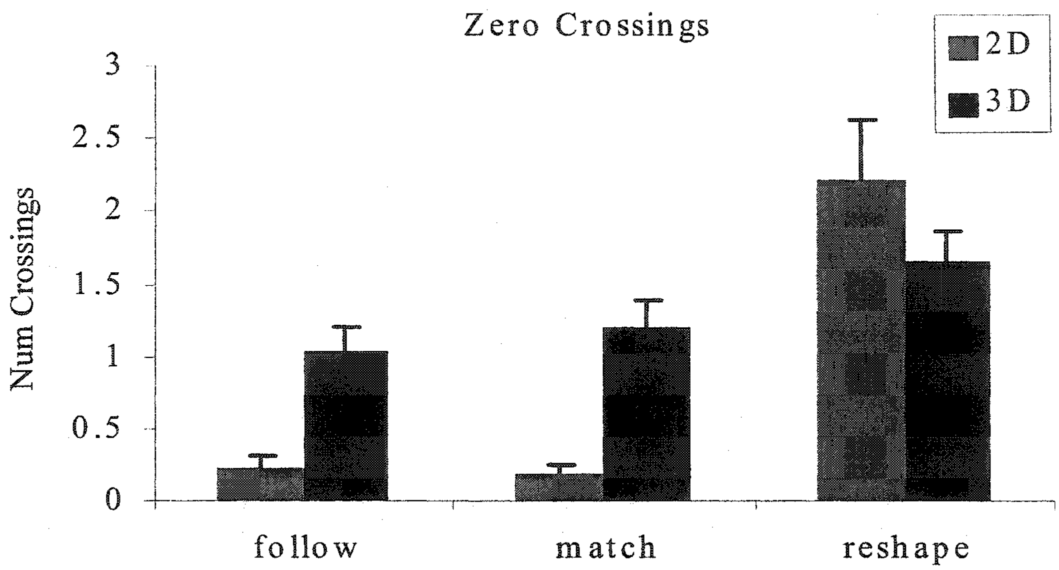

Figure 4.8 - The effect of the *shape x task type* interaction on *zero crossing*

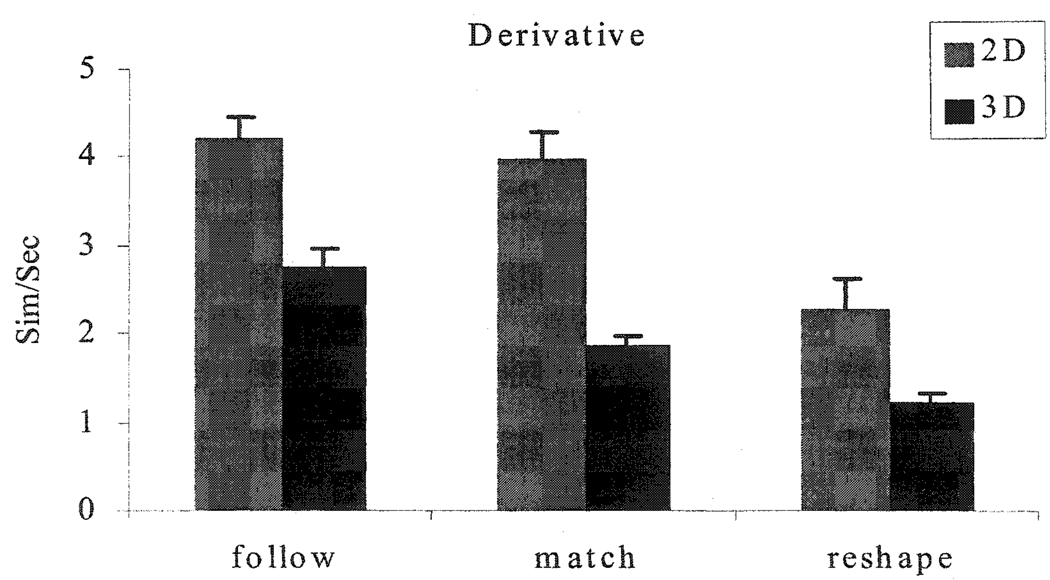

Figure 4.9: The effect of the *shape x task type* interaction on *derivative*

4.8 Test Comparison Study

4.8.1 General

Having studied the sensitivity of Cognitive Cubes to factors related to cognitive performance, we turn to a comparison of Cognitive Cubes to a known tool for 3D spatial assessment: the Mental Rotation Test (MRT) [112,128,167]. As we discussed in Section 3.3.2 the MRT has 3D and spatial components, like Cognitive Cubes, leading us to expect a strong relationship between the two assessments, particularly with 3D tasks. However, since the MRT does not include any of Cognitive Cubes' constructional, planning, or motor task components, we might anticipate the relationship to be limited to simpler tasks such as *follow.*

4.8.2 Participants and Methods

The test comparison study's 12 participants had ages ranging from 18-36, with an average age of 27.66 and standard deviation of 5.61 years. 4 of the participants were females and 8 were males. Participants were all volunteers recruited on and off campus none of whom participated in any other phases of the Cognitive Cubes experiments. The procedure followed the general methodology except that participants took the MRT test before and after the Cognitive Cubes assessment.

4.8.3 Results

A few qualitative randomly selected results based on post-session interviews with the participants were:

- Male, healthy, 34—"Cognitive Cubes were easier and more fun than the MRT, since they are tactile there's less to think about".
- Female, healthy, 28—"MRT is more challenging than Cognitive Cubes. With Cognitive cubes you can always 'try', with the MRT you have to use only your imagination".
- Female, healthy, 29—"MRT is stressful, feels like an exam. Cognitive Cubes is fun, you feel you are doing something. Less stress and there's no time limit".

Unlike the cognitive sensitivity study, this experiment used a homogenous set of participants without any elderly or AD participants. Completion rates were therefore uniformly high.

Cognitive Cubes results for all four measures are compared to MRT results obtained both before (pre-CC) and after (post-CC) the Cognitive Cubes assessment. Interestingly, post-CC MRTs are markedly improved (almost all in the 90th percentile). While it is well known that repeating the MRT brings improved performance, improvements are in this case well above the normally reported repetition improvement rate of roughly 5%.

We performed our MRT/Cognitive Cubes comparisons using correlations. Because they reached ceiling and lost sensitivity, correlating post-CC MRTs to Cognitive Cubes measures would be a meaningless exercise and are not presented.

The correlations of pre-CC MRT and Cognitive Cubes are presented in Table 4.5, along with the probability that the correlations are not significantly different from 0. Those correlations with high probability of being different from 0 are presented in bold, underlined digits. The measure with the most significant overall correlation (and the only

Table 4.5 – Pre-CC MRT/Cognitive Cubes correlations. Correlation significance: (p<.1) in bold, (p<.05) underlined

Dependent Measures	**Overall**	**Shape Type**		**Task Type**		
		2D	**3D**	**follow**	**match**	**reshape**
Last connect	-0.38	**-0.49**	-0.35	**<u>-0.63</u>**	-0.35	-0.24
Similarity	0.03	-0.36	0.17	0.16	-0.09	0.08
Zero crossings	-0.23	0.07	-0.25	-0.14	-0.45	0.11
Derivative	**0.51**	0.38	**<u>0.57</u>**	0.43	**0.50**	0.34

reaching marginal significance) is the *derivative* measure. Correlations to *zero crossings* are low. Correlations to *similarity* are also low, perhaps because similarities are uniformly high. Correlations to *last connect* are also high. Correlations are only slightly stronger for 3D than 2D shapes, while correlations are strongest with *follow* tasks, slightly weaker with *match* tasks, and completely untrustworthy with *reshape* tasks.

4.9 Discussion

4.9.1 Limitations

We believe there are several reasons for caution when drawing inferences from our experimental work. First, our experiments involved a limited number of participants and were motivated by the desire to improve and learn about the Cognitive Cubes tool, and not by basic scientific questions about cognitive function. This led us to choose an experimental design that was not balanced in *shape type* or *task type*, so that we could emphasize those tasks that seemed to us most promising in assessment. ANOVA results should therefore be interpreted with care.

We did not randomize or counterbalance ordering of *task type* and difficulty, instead we used a rough order of increasing difficulty. This enabled participants beginning to struggle with the current task to skip following tasks with which they most likely also would struggle, which proved crucial in retaining elderly and AD participation. At the same time, this decision not to randomize or counterbalance introduced a practice effect that must be reckoned with in the effect of the *task* type factor.

Since the elderly and AD participants struggled with tasks most often, only the stronger of these participants completed the more difficult tasks (Table 4.2). This unbalanced the age factor and made task performance by these participant groups appear better than it would otherwise. Even so, mean task performance by elderly and AD groups was still worse than performance by the young.

Finally, we introduced filtering and repetition into our trials to handle the remaining hardware shortcomings of our prototype. The frequency of repetition was relatively low, and analyses excluding them were very similar to those shown here. Filtering was more frequent, and since structural collapses were more likely with elderly participants and during 3D tasks, it may have distorted the results of the *age* and *shape type* analyses. However, since our filter only affected Cognitive Cubes events which removed and immediately replaced multiple blocks in less than ten seconds (otherwise an infrequent and discouraged participant action), we are confident that we have controlled this potential problem well.

4.9.2 Confirmations

First, we note that the ActiveCube hardware component of Cognitive Cubes performs quite well, given it is a prototype. While spot repairs were sometimes required, the hardware continued to function through well over 50 hours of constant, demanding use. Participants had few complaints, were engaged and interested, and were usually having fun.

Spatial cognitive performance is known to decline with increasing age, cognitive load, and shape complexity, so it is reassuring and gratifying to see these trends in Cognitive Cubes' measures. The only exception to this trend is in the effect of *task type* shown by *similarity*, it is lower for the *follow* task, and thus less similar to the target despite less cognitive load. Since *follow* was always the first task type, this may be a side effect of ordering: participants had not yet reached peak performance when they were performing *follow* tasks.

Preliminary results indicate that Cognitive Cubes is sensitive to mild AD. Though we examine Cognitive Cubes with only two mild AD participants (Figure 4.5 and Table 4.4), outcomes show strong differences between unaffected elderly and AD participants. Further work should examine whether Cognitive Cubes can discriminate between AD and other explanations of constructional weakness.

Finally, not only does Cognitive Cubes respond well to known cognitive factors, but certain of its component measures also have sensitivities similar to a 3D assessment already in wide use: the MRT. Other Cognitive Cubes components promise additional sensitivities not available in the MRT.

4.9.3 Surprises

Since our primary goal was system evaluation and not cognitive research, we did not form any hypotheses about interactions among the cognitive experimental factors of *age*, *task* type, and *shape* type. However, we were pleasantly surprised that Cognitive Cubes is sensitive to an interaction between *task* and *shape type*. One possible explanation of the interaction is that with 2D target shapes, the additional cognitive planning load of *match* (vs. *follow*) is minimized. At the same time, since the starting point in *reshape* is 3D, 2D targets still require significant cognitive planning. With 3D target shapes (by all measures except *last connect*), the cognitive load increases steadily from *follow* to *match* to *reshape*. The *last connect* exception may indicate the added time it takes to move from a 3D to a 2D target shape. Alternatively, it may result from the combined effect of participant dropout and practice.

Contrary to our expectations, both 2D and 3D task types produce some good correlations to the 3D MRT. We believe this may well be attributable to task difficulty. While the MRT asks the user to perform a small set of relatively simple 3D mental rotations, Cognitive Cubes challenges participants to construct a single shape, which may be small or large, 2D or 3D. Which is more like the MRT: building from scratch a complex 3D shape, or a simple 2D shape? The answer is unclear, and thus the lack of clarity in the *shape type* correlations.

The improvement from pre- to post-CC MRT is unexpected, but very intriguing. Could Cognitive Cubes be used as a form of cognitive therapy or training, for example in rehabilitation? (see also Section 6.3).

4.9.4 What's Next?

Is Cognitive Cubes a useful tool for assessment? Our experience certainly indicates great promise. Despite being a prototype, the ActiveCube hardware component stood up well to intense use and proved to be quite intuitive for our participants. In experimental evaluation, the system as a whole was sensitive to well-known cognitive factors and compared favorably to an existing assessment. Automation introduced a previously unachievable level of reliability and resolution in 3D measurement and scoring. Despite all this, Cognitive Cubes is not yet ready for regular use.

How might Cognitive Cubes be prepared for use in the field? The gap between a good prototype and a reliable tool is a large one. Use in clinical or research settings would require significant improvements in cost, reduction of connection and system errors, and improvements in structural strength. These are fairly typical requirements for the development of any technology. In addition, extensive testing would be required to identify the distribution of scores typically achieved with Cognitive Cubes. In this way, assessors can reliably decide whether or not a score is exceptional (see also Section 6.3).

How might Cognitive Cubes be improved? The system could be greatly improved with a more polished notion of task difficulty, which then might be used to weigh assessment results over multiple tasks into a composite score. In this study we use *shape type* as a rough approximation of difficulty (without weighting composite scores), but certainly the number of blocks needed to build a shape should also be a factor. Researchers and thinkers in a variety of fields have proposed numeric measures of shape complexity [82]; Cognitive Cubes could be a good mechanism for testing their relevance to humans.

One the most unique strengths of Cognitive Cubes is its ability to capture each step of the task progress—closely mirroring the cognitive processing of the participant. With the same data used to build similarity graphs it is also possible to build decision trees reflecting the participant's chosen path through the space of possible cube-by-cube construction sequences. This dynamic process can be probed even more deeply by attempting to categorize participant trees according to cognitive ability (see also Section 6.3).

4.10 Conclusion

In this chapter, we described the design and evaluation of Cognitive Cubes, to our knowledge the first system for the automated assessment of 3D spatial and constructional ability. Cognitive Cubes makes use of ActiveCube, a 3D spatial TUI, for describing 3D shape. Cognitive Cubes offers improved sensitivity and reliability in assessment of cognitive ability and ultimately, reduced cost. Our experimental evaluation with 43 participants confirms the sensitivity and reliability of the system.

We see Cognitive Cubes as a proof-of-concept demonstrating our research goal, showing that a specialized spatial TUI closely tied to an application can offer substantial benefits over existing solutions and suggests completely new methodologies for approaching the applied problem.

We also see Cognitive Cubes as a practical and successful example of our TUI heuristics being put to work. Our choices during the design of

Cognitive Cubes, for example, selecting a spatial application domain, and choosing a very intuitive spatial mapping between the interface and the task, were all closely guided by our TUI heuristics. We believe that the success of Cognitive Cubes should also be attributed to the heuristics that guided the design.

Chapter V
The Cognitive Map Probe

5.1 Introduction

The second half of our proof-of-concept is the Cognitive Map Probe (CMP). Using the CMP (Figure 5.1) we tried to demonstrate again that we can solve a real-life problem using spatial TUIs, closely following our design heuristics, providing benefits over existing solutions and

Figure 5.1 – An Overview of the Cognitive Map Probe

suggesting new approaches that were not conceivable without a TUI.

As we did in Cognitive Cubes, we approached the cognitive assessment domain with a specialized TUI, this time looking at the assessment of cognitive mapping abilities. The idea behind the CMP is quite simple, and follows the spirit of Cognitive Cubes: show participants a spatial prototype, ask them to memorize it and later ask them to reconstruct it with a spatial TUI.

Why did the CMP deserve its unique part in our proof-of-concept? Like Cognitive Cubes, the CMP uses a physical, spatial TUI, and a virtual display. However, the CMP differs greatly from Cognitive Cubes. The CMP is attempting to offer a high level of task realism, and in this sense, a high level ecological validity, unlike the more abstract, mind-puzzle approach of Cognitive Cubes. The CMP's virtual prototype is a neighborhood, presented from a first-person perspective. For its physical interface, the CMP uses a set of realistic looking miniature-models of buildings. However, unlike Cognitive Cube, the CMP interface is essentially 2D in nature, following a tabletop interaction paradigm.

Like Cognitive Cubes (see Section 4.1), the CMP design adhered our spatial TUI heuristics (See Section 2.2). The TUI offers a highly intuitive spatial mapping (Section 2.2.2) to the assessment task. The physical building models allow natural placement on the CMP board. The physical buildings also have a very close visual resemblance to their virtual parallels—they are practically the physical embodiments of the models, as they were printed from them using a 3D printer (Figure 5.1 and Figures 5.6-5.8).

As with Cognitive Cubes, we argue that the CMP does offer a strong I/O unification (Section 2.2.3). In the CMP, like in Cognitive

Cubes, the display and the prototype presented on it are separated from the TUI. However, in the cognitive assessment setting the prototype is external to the interaction. Moreover, in the CMP (unlike in Cognitive Cubes) the prototype display is turned off as soon as the interaction with the TUI begins. The participant interacts solely with the TUI, which naturally integrates input and output—actions performed on the physical models register immediately and fully (building identity, position, orientation and time of action) at the host computer.

Like Cognitive Cubes, and like most commercial construction sets, the CMP fully supports trial-and-error actions (Section 2.2.4). The participant can place or remove objects from the TUI in any desired order. However, unlike Cognitive Cubes, in the CMP this behavior is restricted to a 2D, tabletop interaction.

The CMP is attempting to measure a somewhat elusive cognitive ability. As we detailed in Section 3.2, wayfinding is an essential life skill, relying on a person's cognitive mapping ability, or the ability to construct mental representations of an environment. Age, disease or injury can severely affect cognitive mapping, making assessment of this basic survival skill particularly important to clinicians and therapists. As we detailed in Section 3.2, many techniques have been developed over the years for measuring and assessing cognitive mapping ability. Map drawing or placement is quite common, but is difficult to score consistently, wholly two-dimensional and necessarily quite abstract in representation. A few researchers have assessed cognitive mapping by asking patients or study participants to arrange 3D objects representing elements of their environment. While the reduced level of abstraction and more 3D representation likely increases assessment sensitivity, previous implementations of this approach were quite unwieldy and still inconsistently scored.

The CMP was designed to address these problems in assessment, being an automated tool for the measurement of cognitive mapping ability. The CMP makes use of the Segal model [52,53], a tabletop TUI originally designed for the input of architectural models. CMP users view a drivethrough of a neighborhood on a large screen perspective display, and then input their recollection of that neighborhood by arranging 3D building models on the Segal model's tabletop input surface. The CMP automatically records and scores each change the user makes to the model configuration.

The Cognitive Map Probe is the first TUI for the assessment of cognitive mapping ability, combining intuitive spatial mapping, natural I/O unification, support of trial-and-error actions, and increased ecological validity and sensitivity of 3D input with the improved consistency, efficiency, flexibility and high-resolution data collection of computerization.

In this chapter we detail the hardware and software of the CMP and present a rigorous experimental examination of the sensitivity of the CMP to age and task difficulty, two factors that have a well-known relationship to cognitive mapping performance.

5.2 Hardware

5.2.1 Infrastructure – the Segal Model

The Segal model (see Figures 5.1, 5.2 and 5.8) is a pioneering TUI named in memory of architect and advocate of home self-design Walter

Segal. In 1982 John Frazer and his colleagues built the Segal model to support Segal's work on self-design of houses [52,53]. Segal had developed a timber-frame technique enabling home-building without professional help, but found that users of his technique encountered major difficulties when visualizing their homes or designing from a blueprint. In order to help users with no experience or knowledge of architecture or computers Frazer and his colleagues designed the Self-Builder Model for Segal (the model was named the "Segal Model" after Segal died).

The device was designed to enable tangible, direct interaction with architectural floor plans and their components, such as walls, doors, windows, plumbing fixtures and furniture. It is a tabletop size (102cm × 71cm) board covered with an array of 768 edge connector slots arranged in 24 columns of 16 vertical slots and 16 rows of 24 horizontal slots. Each slot has contacts enabling recognition of 120 different entities, after accounting for symmetries in orientation. Architectural components were represented by simple physical 3D models, with each type of component coupled to a unique connector identification code.

The original entities were colored plastic panels (Figure 5.2), representing walls and windows, along with several other small-scale wooden objects, representing furniture and other home-entities. When the user places the physical objects on top of the model, the position and identification of these objects is sampled electronically in real-time. In the original design domain the output was used as input to a wire-frame rendering software package. The software included a design feedback tool which synthesized "expert advice", such as house area and cost, to interactively help the user (and in a way, imitate some of the expertise of an experienced architect [52]). The board was scanned electronically, in real-time with very low processing demands from the hosting system (that is, no real-time image processing, etc.).

As was discussed in detail in Section 2.8, there are several other modern tabletop TUIs. Our choice of the Segal model as the infrastructure for the CMP was based on one major advantage: accessibility. Tabletop TUIs are still one-of-a-kind prototypes which are generally being used by their developers. John Frazer generously offered the Segal model to us for our CMP research, warning us that the model was in storage and hadn't been used for many years. Using the Segal model enabled us to focus most of our efforts on designing, implementing and testing our cognitive assessment application, the Cognitive Map Probe, rather than on the lengthy engineering effort of developing a new tabletop TUI. In retrospect, as will be detailed later in this section, the Segal model was a well suited choice for the CMP infrastructure and enabled us to address all the research questions we were planning to pursue.

5.2.2 Reviving the Segal Model

The Segal model is a one-of-a kind device, a single copy was constructed for the research, and there are no known copies. The model was used with an early 1980's, now obsolete, computer [52,53,154]. Our initial goal was to "make the Segal model work" with a standard, modern PC using an interface that was as generic as possible. We have revised the Segal model hardware and software interface to enable the model to be scanned using the parallel port of a standard PC. The new interface is able to read out the entire board at a sample rate of about 500Hz (for detailed discussion of our efforts in revising the Segal model, please see [154]).

A simple, preliminary demonstration of the power of the Segal model as a design tool was achieved when we rendered a tangibly constructed world via a "Half-Life" [147] computer-game 3D graphics engine. For this preliminary effort we used John Frazer's original physical objects as interfaces, namely the set of colored plastic panels as representations of walls in the "Half-Life" world.

The user started by building a world using physical tangible objects on top of the Segal model base (the user is physically building the letters "U OF A", see Figure 5.2). The physical model is sampled and transferred automatically to a virtual model in a "Half-Life" based 3D environment, with full control of the appearance of the world (texture, lighting, etc.) and full ability to take a virtual walkthrough in the former physical model (Figure 5.3). Furthermore, the virtual world can now be populated with dynamic virtual entities and characters, either by physical-tangible means or by software editing (Figure 5.4). The physical model can then become a fully interactive, fully active virtual world (Figure 5.5). For further details, see [138,140-142,154].

5.2.3 Printing 3D Interfaces

Our design goal was for the CMP to be a TUI for inputting cognitive maps in assessment tasks. We envisioned an interface that would allow us to introduce an adult user to a detailed virtual environment, enable cognitive mapping, and then assess the perceived cognitive map using a TUI (the CMP). Thinking of the elderly and of people suffering from mild AD we wanted the virtual environment to have the appearance of a

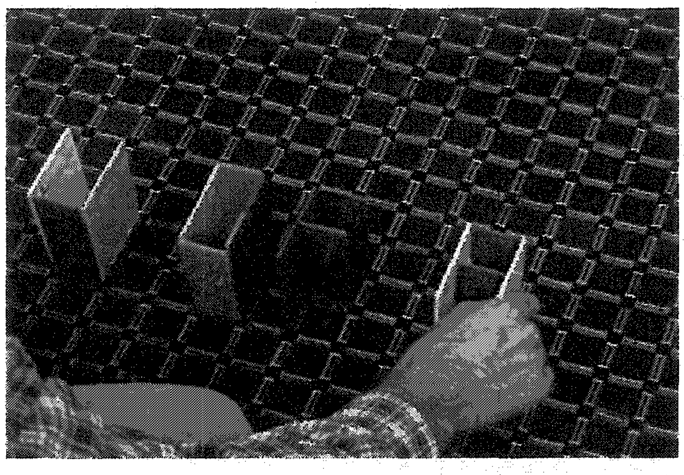

Figure 5.2 - Building a world tangibly

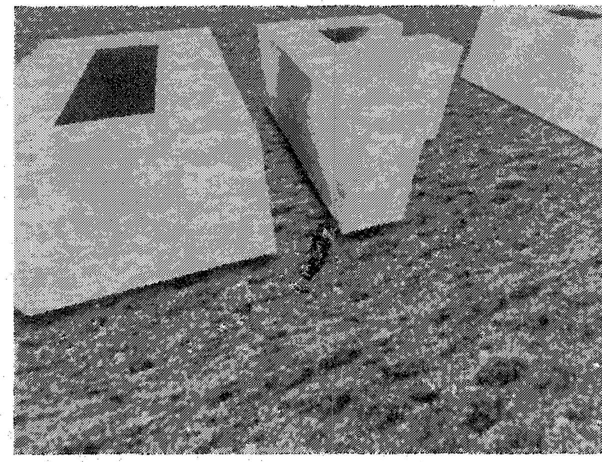

Figure 5.3 - The virtual "OF" structure

Figure 5.4 - Avatars inside the letter "U"

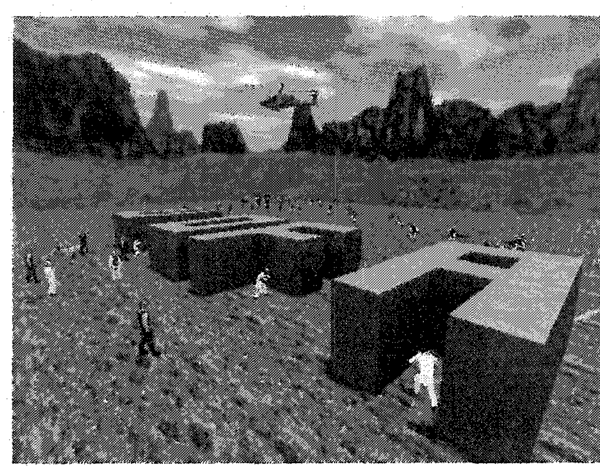

Figure 5.5 – Virtual battle scene

realistic neighborhood. We knew that the subjects for our work would all be from the local community, so we wanted the neighborhood to appear as similar as possible to a typical Edmonton neighborhood [90]. To achieve these goals, the CMP hardware had to:

1. Describe a compelling 3D editable neighborhood, using small scale models of buildings as the environment entities and interfaces.
2. Maintain highly intuitive spatial mapping to the assessment task. For this, the buildings had to be realistic, detailed and

convincing. The buildings had to fit well into a typical Edmonton residential neighborhood, since those were the neighborhoods our participants would know. The appearance of the interface should appeal to adult users, and not have the look and feel of a simplified toy.

3. Maintain extremely high resemblance between the hardware interfaces (the physical buildings) and their virtual counterparts which are displayed as part of the virtual environment.

With these goals in mind we considered and rejected several possibilities for the design of our interfaces (for example, using train set houses, construction sets, etc. [153]). With the help of Prof. Robert Lederer, his student Adrien Cho and their colleagues (from the University of Alberta Industrial Design department), we chose a new approach that answered our design goals – we decided to print our interfaces.

We chose several real architectural landmarks that are a familiar part of many Edmonton residential neighborhoods. These included several different residential houses, a church, a fire station, a strip mall, a supermarket and a gas station. All buildings were photographed and then modeled, with the photographs as guidelines only, using a 3D CAD (Computer Aided Design) tool (Rhinoceros©). We ended up with 10 highly detailed virtual 3D building models. These 3D virtual models were printed using rapid prototyping technology (a Genisys 3D printer). The resulting polyester objects are quite sturdy and mounted on flat bases, under which is a single connector for the Segal model's board (Figure 5.6).

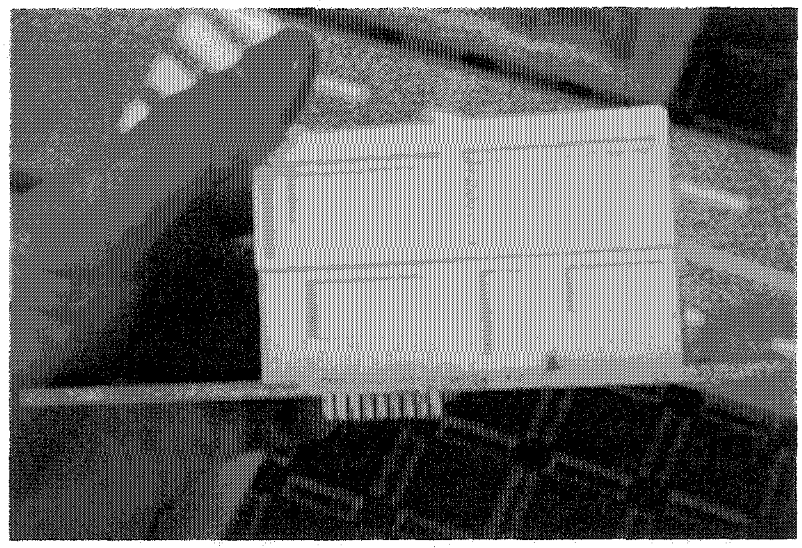

Figure 5.6 – Printed house model with its edge connector

Aligning the base with the board's slots aligns the connector to its matching slot and eases insertion of the model (Figure 5.7). All the models are of similar scale and can be arranged easily with two hands. The user can connect the buildings in any desired orthogonal orientation (by using either horizontal or vertical slots on the board and inserting the building in one of the two possible symmetrical aspects). The models were spray painted in single primary colors for easy viewing by the

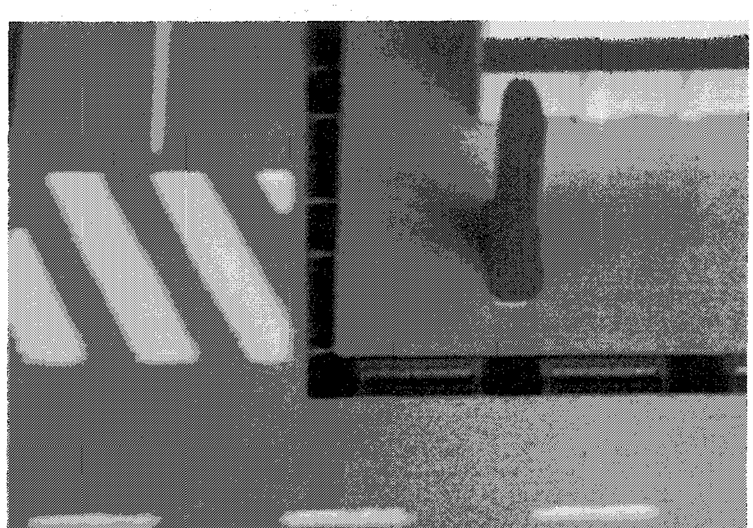

Figure 5.7 – Alignment of the model base and the CMP slots

elderly (for a discussion of visual affects of age in elderly people see [36,132]), but important details such as store signs were hand painted in contrasting colors. The models are quite detailed in shape, and include doors, windows, and even the patterns of wood siding.

We also attached a simple street pattern to the board (one four-way and one "T" intersection; see Figure 5.8); this street pattern was never removed during assessment. The street pattern was also designed using a 3D CAD tool but due to its size was not printed but rather cut out of a plastic sheet using a Computer Numerical Controlled (CNC) machine. All 10 models and the street pattern can fit onto the board at the same time. A 3D compass model, presenting an arrow pointing to the environment's

Figure 5.8 – A participant interacting with the CMP

north was also designed and printed. An experimental administration tangible control object was also used as part of the CMP hardware. The object was connected to (or disconnected from) the CMP by the assessor in order to start, end, or advance the experiment to the next neighborhood (for further details on the design of the models see [21]).

5.2.4 The CMP *Neighborhood – Physical vs. Virtual*

Practically, all the physical models we created for the CMP are TUIs: they can be manipulated and positioned by the user on top of the CMP, while their position, orientation and identification are being sampled. Our design approach enabled us to effectively "print out" interfaces, keeping their form identical in both the virtual and the physical domains.

Two minor mismatches between the virtual and physical neighborhoods do exist. We decided against adding these device characteristics to the virtual neighborhood because we felt it might reduce its perceived realism.

1. The virtual environment does not present the neighborhood as containing edge connector slots (which track the buildings position and identification), and in this sense contains a mismatch to the physical neighborhood.
2. The physical plastic bases which the buildings are mounted on (for alignment purposes) do not appear in the virtual neighborhood.

5.3 Software

5.3.1 Facilitating Cognitive Mapping

The CMP supported cognitive mapping of the neighborhood using an interactive 3D virtual environment (Figures 5.9 and 5.10). The virtual environment was implemented using the SVE VR toolkit [74] and enabled viewing and interaction with any given neighborhood layout. The 3D environment was projected using a monoscopic digital rear projector on a large screen (205cm diagonal). The interactive 3D virtual environment allowed us the choice of numerous interaction paradigms with the neighborhood model. Viewing could be passive, similar to riding a bus as a passenger, or active, similar to driving a car. Viewing could be

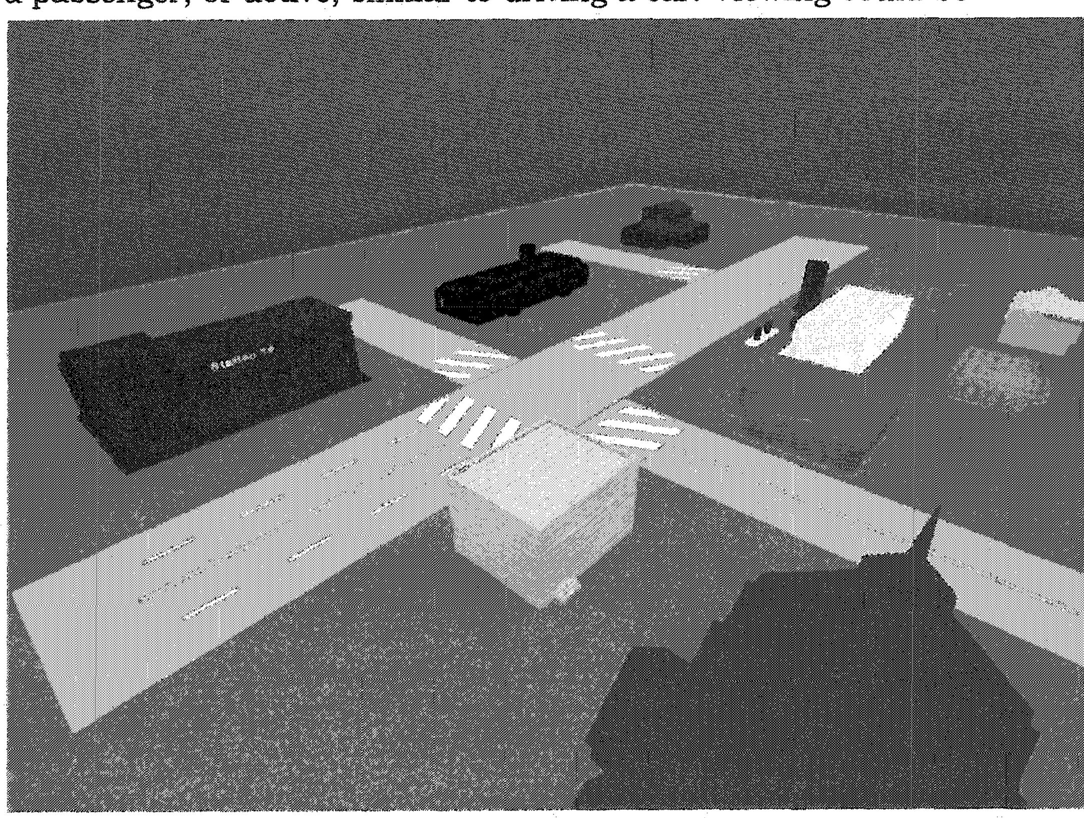

Figure 5.9 - Virtual neighborhood (exocentric view shown for illustration)

egocentric, with participants viewing a street level view of the neighborhood. Viewing could also be exocentric, with participants seeing a bird's eye view. It would also be easy to change other parameters of the virtual neighborhood, including different weather conditions (snow or rain) and different lighting conditions (day or night) [178].

When choosing the interaction technique with the 3D virtual neighborhood we had to take into account our subjects' age and abilities. Implementing a demanding interaction technique for the mapping phase might not allow some of the subjects to properly learn the new environment. Some elderly subjects might not be able to perform well with the CMP simply because they were not able to learn the new environment when using a novel walkthrough interface. On the other hand, a completely passive drivethrough might hinder the development of the subject's cognitive map [83].

We decided on the following interaction metaphor for the virtual walkthrough of the neighborhood (see also Figure 5.10):

1. The participant is a passenger in a bus driving through the virtual neighborhood. The experiment assessor is the bus driver.
2. The bus moves through the neighborhood along the same path, covering all of the neighborhood streets in a systematic order (taking a right turn at all junctions, and a U turn at all dead ends, beginning and ending the ride at the same point in the south of the neighborhood).
3. The participant views the virtual neighborhood from a passive, egocentric perspective, moving through the neighborhood at street level.

Figure 5.10 – The CMP drivethrough – participant's view.

4. The participant is interacting with the bus driver, asking for the bus ride to start.
5. At any point during the bus ride the participant can ask the bus driver to stop the drivethrough to allow further examination of the neighborhood. After this request the drivethrough will resume only when the participant asks for it.
6. At any point the participant can ask the bus driver to rotate slowly through 360 degrees for a panoramic viewing before continuing along the viewing path.
7. A virtual compass in the ground plane indicates which direction is north to help the participant orient herself before the drivethrough begins. It is displayed at the starting point (the

south of the neighborhood) and disappears when the bus starts moving.

8. A semitransparent sign welcomes the participant to the drivethrough and notifies her of the end of the ride. The welcoming sign disappears as soon as the drivethrough begins.
9. In order to establish a consistent drivethrough experience, all interaction events, including the drivethrough the pre-defined path, are fully automated. The assessor only has to press a button in order to start, stop or rotate the bus, following the participant's request.

5.3.2 Probing

During the experiment, the CMP collects the following data:

1. Drivethrough data: time, action (ride start, ride end, ride stop, ride continue, ride rotation), and position (in which the action took place).
2. Cognitive map data: a vector of events; each event contains the following fields: time, object ID (buildings ID or control object), action (connection/disconnection), position and orientation.

The main measures used in our analysis were based on the cognitive map data. After the assessment session, the CMP analyzes the collected data to score the participant's performance. As discussed in Section 3.2, there are a number of ways in which cognitive maps may be scored. All of these methods involve comparisons of the actual map *M* to the participant's cognitive map *C*.

The measures that are implemented as part of the CMP can be divided into four groups: (1) Set measures (*number* and *difference*), (2) Intersection measures (*distance, orient,* and *interbuilding*), (3) Similarity measure (*similarity*) and (4) Overall measures (*totaltime* and *dSim*).

Measures that disregard position and treat M and C only as sets of buildings are (note that we use the annotation $|X|$ for the number of members in the set X):

$$(5.1)\ number = 1 - \text{abs}(|M| - |C|)/|M|$$
$$(5.2)\ difference = 1 - (|M\text{-}C| + |C\text{-}M|)/(|M| + |C|)$$

Measures that compare position only within the set of intersecting buildings $M \cap C$ include:

$$(5.3)\ distance = 1 - \Sigma_i\ (dist(M_i, C_i)/\mathrm{d}_{max})/|M \cap C|$$
$$(5.4)\ orient = 1 - \Sigma_i\ (odiff(M_i, C_i)/180)/|M \cap C|$$
$$(5.5)\ interbuilding = 1 - \Sigma_i\Sigma_j\ (\text{abs}(\mathrm{D}_{Mij} - \mathrm{D}_{Cij})/\mathrm{d}_{max})/|M \cap C|^2$$

where all sums range over the set $M \cap C$, *dist* is Euclidian distance, *odiff* is the angular difference in degrees between the orientation of two buildings, d_{max} is the length of the CMP board diagonal, $|M \cap C|$ *is the number of members in the set* $M \cap C$, and D_M and D_C are square matrices in which the entries are $dist(M_i, M_j)$ and $dist(C_i, C_j)$, respectively, with i and j ranging over the set $M \cap C$. Finally, the CMP forms a composite measure that includes both set and position error:

$$(5.6)\ similarity = difference \times distance \times orient$$

Recall that the CMP also records the time of each action on the board. This allows us to add *totaltime*, the time it takes to complete one assessment trial, to our suite of measures. We can also probe the progress participants make during the assessment by comparing our measures to the current time. We construct the additional measure *dSim* by finding the differences between consecutive measurements of *similarity* (Equation 5.6) divided by the time elapsed between those measurements, and averaging the resulting "local slopes" over all such pairs in an assessment trial.

5.4 System Strengths and Weaknesses

The CMP has some technical limitations:

- *"Hard" Connectivity.* The interaction with the CMP is based on "hard" physical connection: an electrical connector is inserted into an electrical slot (see Section 5.2.3). Requiring this explicit connecting action from the user can be intrusive. "Soft" connectivity, based on a different flavor of a tabletop TUI (for example, vision based, see Section 2.8) might facilitate easier interaction. However, these other methods have their own shortcomings.
- *Orientation inputting.* The CMP limits the orientation inputting of the models to orthogonal angles only (0°, 90°, 180° and 270°). It follows that the interface limits the orientation of buildings in the prototype neighborhood, and the variety of orientation errors a participant can potentially make. A more flexible orientation input scheme might allow us to add more orientation complexity to the

buildings in the prototype neighborhood, and perhaps enhance the probing resolution of the participant's orientation errors.

Despite these limitations, the CMP is the only computerized, tangible system for assessing cognitive mapping ability.

The CMP offers the following major advantages over existing methods for assessing cognitive mapping skill:

- *Sensitivity*. The CMP monitors participant progress (or lack thereof) throughout map construction. In contrast, existing methods assess cognitive mapping only when the map is complete. In addition, the CMP's 3D tangible interface allows a much more direct translation of cognitive maps into physical representations, with fully detailed buildings viewable in perspective from all sides, much as they are during travel through the represented neighborhoods themselves. Commonly used 2D cognitive mapping assessment methods offer only highly abstracted 2D projections of the represented environment and its buildings (see Section 3.2). Ultimately, it should be possible to add adaptivity to the CMP, focusing more quickly and completely on the limits of participant ability, and improving sensitivity further.
- *Accessibility*. Many of the populations commonly given cognitive mapping assessments face cognitive, visual or motor challenges. Unlike traditional 2D assessment techniques, the CMP uses an interface that is intuitive, easy to see, and simple to manipulate. This proved invaluable during our work with the elderly.

- *Consistency.* If an assessment is to have meaning outside of its original context, it must be performed consistently and reliably by all assessors. Existing 2D assessments are consistent, but achieving this consistency requires that the assessments be fairly simple to perform, reducing assessment sensitivity. Because it is automated, the CMP achieves the highest level of consistency while at the same time improving sensitivity with complex tasks and very frequent measurement of the participant.
- *Control.* The CMP's virtual neighborhood display will always be simpler than real-world stimuli. On the other hand, virtual display offers an amazing degree of control in assessment. Climates can be changed, landmarks rotated or removed, buildings located incorrectly by the participant can be displayed translucently on top of correctly located buildings, and neighborhoods can be viewed from positions in midair – effects extremely difficult if not impossible to achieve in the real world.

5.5 Cognitive Sensitivity Study

5.5.1 Methodology

Is the CMP useful? How sensitive is it to well-known cognitive factors in practice? We explored these questions with an experiment. The CMP was designed to support a wide range of cognitive mapping tasks. In our experiment, we sampled this range by varying the number of buildings in the virtual neighborhood we asked participants to recreate.

We expected that cognitive mapping performance would worsen by all measures as the number of buildings in the mapped environment increased. We also anticipated that performance among our elderly participants would be worse than the performance of our young participants, reflecting the natural effects of age on cognitive mapping ability. The experiment was conducted with the approval of the University of Alberta Health Research Ethics Board (HREB-B, file number B-061201-COMSCI). See Appendix B for the experiment information letter, consent form and detailed protocol.

5.5.2 Participants

Our experiment had 20 participants. They were recruited from on and off campus, and ranged in age from 25 to 81. There were ten young participants under 55 years old, ranging in age from 22 to 50 years old, with an average age of 30.5 and standard deviation of 8.31 years. The ten elderly participants were aged 55 years or older, ranging in age from 55 to 81 years old, with an average age of 68.9 and standard deviation of 10.86 years. Both groups were almost balanced in gender with 5 female and 5 male in the young participants group and 6 female and 4 male in the elderly participants group.

In addition to these 20 participants we worked with five more participants whose results are not included for the following reasons:

- Our first participant, a student in his 20's, took part in a pilot study. As a result of our session with this participant we limited the number and adjusted the order of our tasks. His results are not included in our analysis or discussion.
- As a preliminary study, we also worked with one additional 59 years old male participant who had been diagnosed with very mild AD (a Functional Assessment Staging (FAST) [119] score of 2). This

single participant performed all the CMP tasks. His results were not included in any of our experimental analyses or discussion, except for the one detailed in Section 5.6.2.

- Last, three participants were unable to complete all 10 assessment trials. Their results were not included in our analysis or discussion. Two of these participants were AD patients and were not able to perform even the training tasks. The first of these two AD participants suffered from a more severe form of AD (score of 4 or 5 on the FAST) and the second showed signs of psychological stress which prevented us from performing the experiment. The third participant, an elderly lady, was under time constraints and had to withdraw the experiment after performing only 7 tasks.

5.5.3 Design

Each participant performed 3 practice trials (which we termed *World1 – World3*) and 7 recorded trials (which we called *World4 – World10*). Each trial presented a virtual neighborhood with a different number of buildings. All participants viewed the same virtual neighborhoods in the same order, with the number of buildings in the recorded trials increasing from 2 to 8. Neighborhoods were ordered in this fashion so that thresholds in participant cognitive ability could be quickly identified without subjecting participants to unnecessary confusion or frustration.

We designed the different virtual neighborhoods by using the CMP not as an assessment device but rather as a TUI for input of neighborhood models. Figure 5.11 presents a top view of the CMP physical neighborhood in each of the 10 trials the participants were asked to map and reconstruct (*World1 – World10*).

5.5.4 Setup and Procedure

All experiments were conducted according to a strict written protocol, with a script read out loud to each participant. In the script, the participant was introduced to the CMP, the experiment, and its purpose, and then read an information letter. The participant was told that he or she might stop the experiment at any time, and asked to sign a consent form. The participant was interviewed quickly, answering questions concerning age, education and occupation. Participant anonymity was always preserved (for more details on the experimental protocol, see Appendix B). Accuracy was emphasized over speed in the instructions, with participants asked to be as precise as possible, but reminded that the CMP was recording the speed of their actions. Participants were told that there was no time limit, that they may decide when they had finished each task, but that they should do the best they could in reconstructing each neighborhood.

The assessor guided participants through three initial practice trials to train them in the use of the CMP. All practice trials used simple two-building neighborhoods. In the first trial (*World1,* Figure 5.11) the assessor introduced the CMP board and its models, as well as the "bus passenger" metaphor for the viewing of the virtual neighborhoods. The metaphor was introduced along with the pre-defined drivethrough route using a bus model that the assessor moved by hand through the neighborhood on the CMP board. The assessor then talked participants through a viewing of the virtual neighborhood that corresponded to the map already on the board (*World1*). The assessor made certain that the participant was able to transfer the information between the two domains and understood this virtual-physical correspondence. The assessor also demonstrated that the passive viewing (the "bus ride") might be paused at will and that the passenger could then also ask the driver for a panoramic viewing, which rotated the view point slowly through 360 degrees.

In the second trial (*World2*, Figure 5.11), the assessor introduced board interaction to the participant by asking the participant to identify a slight change to the virtual neighborhood during a new virtual tour, that is, the church was rotated 180° and moved to the north side of the main avenue. The assessor then turned off the virtual neighborhood display and asked the participant to adjust the CMP board to match this changed virtual neighborhood, also enabling the participant to practice the physical interaction with the board and the models. In the third trial (*World3*, Figure 5.11), the assessor confirmed that participants completely understood typical interaction by removing all physical models off the CMP, having participants view a completely new virtual neighborhood, and asking them to recreate it on the CMP board, again after the virtual neighborhood display was turned off.

After successfully completing the practice session, the participant began the recorded assessment. This included 7 trials (*World4-World10*, Figure 5.11). Each trial included two phases: mapping and probing. During the mapping phase participants viewed the virtual neighborhood only once. The viewing was from a passive, egocentric perspective, moving through the neighborhood at street level. The virtual compass in the ground plane indicated which direction was north (Figure 5.10).

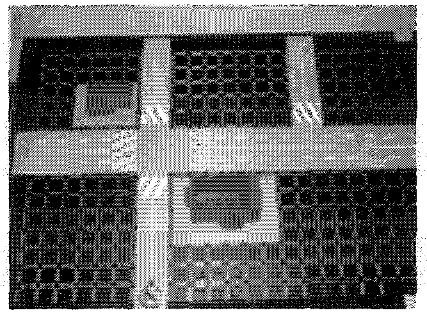

World1 (practice trial)

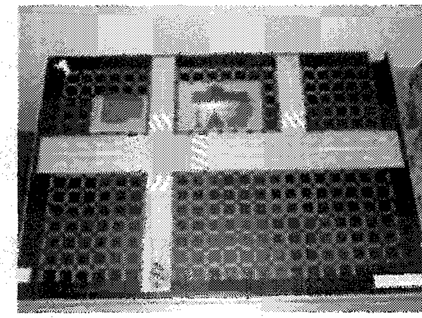

World2 (practice trial)

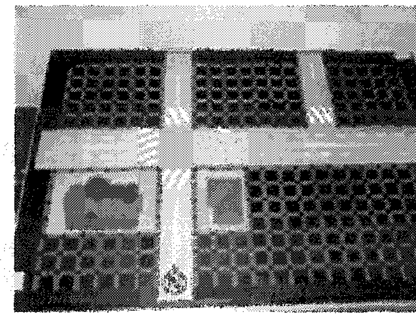

World3 (practice trial)

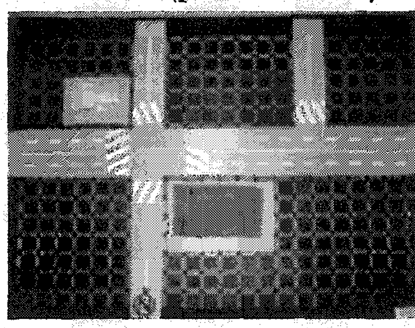

World4 (recorded trial)

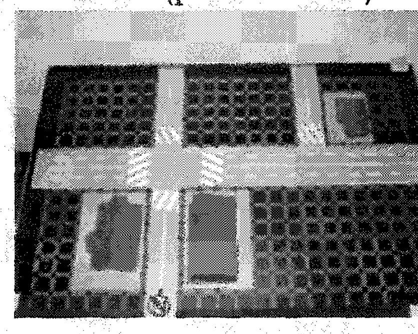

World5 (recorded trial)

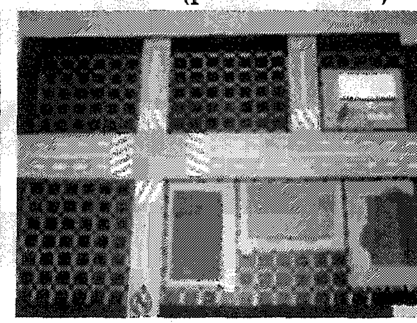

World6 (recorded trial)

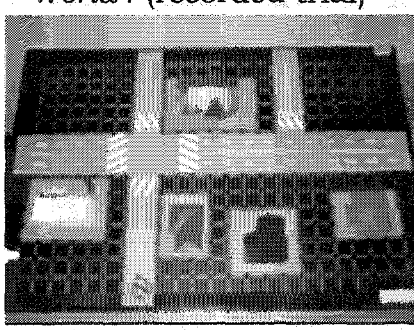

World7 (recorded trial)

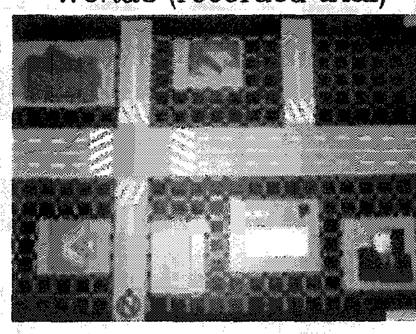

World8 (recorded trial)

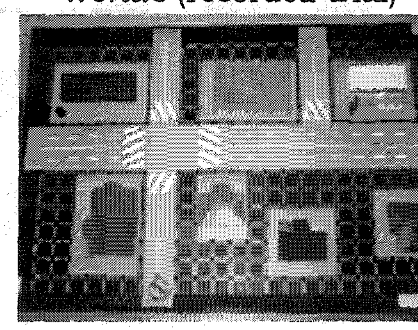

World9 (recorded trial)

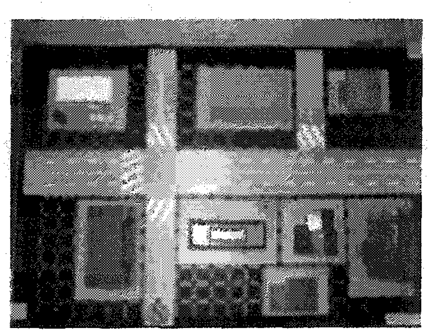

World10 (recorded trial)

Figure 5.11 – The CMP prototype neighborhoods; *World1-World10*

Participants were moved along the same path in each trial. Again, participants could halt their motion at will and rotate slowly through 360 degrees at any time. After finishing the mapping phase the participant moved to the probing phase. The virtual neighborhood display was turned off and participants were asked to interact with the board and to attempt to reconstruct from memory the neighborhood they had just viewed. A physical compass model similar to the compass seen in the mapping phase indicated which direction was north. Participants never received any feedback or comments about their performance from the CMP or the assessor. Participants required 1½ hours on average to complete the full set of 3 practice and 7 recorded trials, as well as a short post-assessment interview.

5.6 Results

5.6.1 CMP—*Measured Reliability*

The 20 CMP participants, each challenged with 7 recorded assessment tasks, performed overall more than 1200 recorded CMP actions (connections and disconnections of objects) overall. Keeping in mind that the Segal model is a historic interface, we fully expected some noise in data collection. However, the CMP performed relatively well (see Section 4.6.4 for comparison to Cognitive Cubes). Most importantly, no participant was forced to repeat a trial. The CMP also made no errors when reporting location. However, there were errors when reporting the identity of the buildings attached or detached from the board. The only such errors that could not be corrected automatically were unidentified buildings, and misidentified buildings.

Unidentified buildings made up 18% of all actions on the board and were corrected interactively by the assessor during the trial. Misidentified buildings made up less than 2% of all actions (21 actions total), but had to be corrected after assessment by manually matching CMP data to video recordings of the assessment. Both error types seem to be correlated with the subjects' skillfulness in physically interacting with the board, for example, subjects who tend to have more difficulty in physically connecting buildings to the board caused more CMP errors. Another more controllable cause of errors was the accumulation of a thin dirt layer on the models' connectors.

Though annoying, both types of errors occurred at rates quite manageable for our purposes and we are confident that a more polished implementation, possibly using different tabletop TUI technology, could eliminate most if not all of these problems.

5.6.2 Qualitative Results and Single Task Example

We start by mentioning a few randomly selected qualitative results, based on the post-session interview with the participants:

- Male, healthy, 62—"Very interesting, encourages thinking, very useful to improve the memory, good training".
- Male, healthy, 59—"Quite demanding, can help improve memory, used colors to remember the houses, was sometimes confused by the 360 rotation in the drivethrough".
- Female, healthy, 79—"Interesting, enjoyed it—fun. The last task was too difficult" (the participant found it hard to believe anyone could remember it).

- Female, healthy, 25—"Driving as passenger was too passive, the ride took too much time".
- Male, healthy, 33—"Not easy, had problems forgetting past drivethroughs (that is, problems 'resetting' memory), used ecological validity when it comes to orientation—front of houses face the street".
- Male, mild AD, 59—"quite liked it, the experiment is a bit too long".

For a somewhat informal view of the results let us examine Figure 5.12. The figure presents the composite measure of *similarity* versus *time* for a single task (see definitions in Section 5.3.2). The *similarity* (Equation 5.6) curves are plotted for 21 participants—10 curves for the young group, 10 curves for the elderly group, and one curve for the single participant with mild AD. All participants were attempting to perform the *World10* task, which had the most buildings of any neighborhood (see Figure 5.11).

Although it is not obvious from the figure, note that none of the participants actually reached a final similarity of 1 (a completely identical match) in the *World10* task. Even the best performers only attained final similarities of 0.97-0.98, accumulating some small distance and/or orientation error.

While most of the young participants were nearly optimal performers as measured by *similarity*, all of the elderly and the single AD participant were not. These older participants attained much lower final *similarity* scores, forgetting or misidentifying one or more of the buildings in the viewed neighborhood.

The *total time* to complete *World10* also varies considerably between the groups. Most of the young participants completed the task

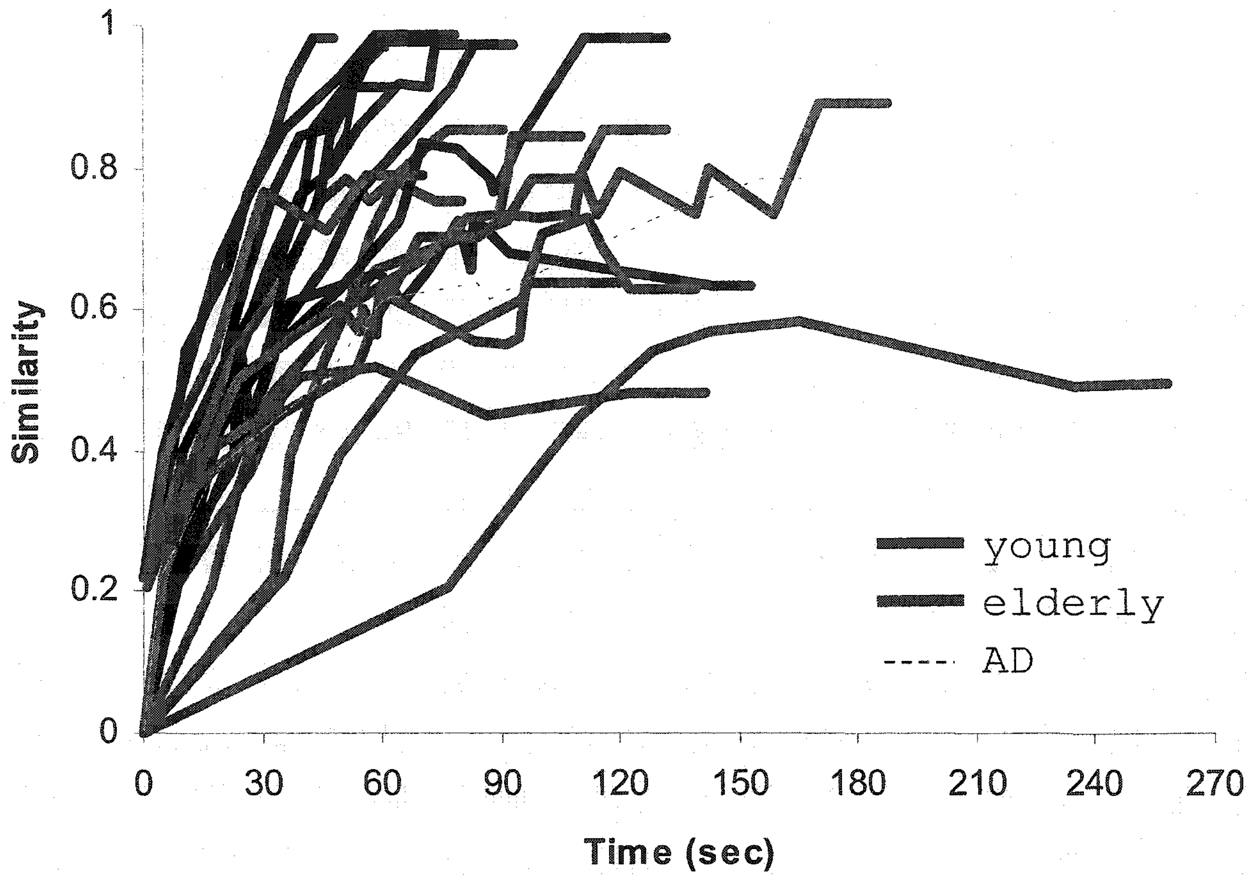

Figure 5.12 – Similarity vs. time; *World10*; all participants

more quickly than the elderly participants, and most of the elderly participants completed the task more quickly than the AD participant.

The slopes of these curves (or their gradients) are directly related to the participants' rate of progress. In *World10*, most of the young participants had curves with much steeper slope than the elderly group, whose curves usually have steeper slope than the AD participant's curve.

Changes in the up/down direction of the slopes (zero-crossing of slope derivative) can be related to participant uncertainty. Changes in slope direction are due primarily to participants adding the wrong building, or removing a correctly placed building. In *World10*, *most* of the young participants had curves with few if any changes in slope direction, while most elderly participants had many changes.

5.6.3 Analyses

Figure 5.13 presents our experimental results by all dependent measures for all 20 CMP young and elderly healthy participants (the single AD participant was excluded from the analysis). We analyzed these results with one ANOVA for each dependent measure. Each such analysis was two-way (2 *age* x 7 *num buildings*), with *age* a between subjects factor, and *num buildings* a within subjects factor. Results of these analyses are presented in Table 5.1.

The CMP responded very much in line with our expectations to the cognitive factor *age* and the task factor *num buildings*. All measures responded significantly to *age*, with the elderly uniformly worse in cognitive mapping performance. In the seven measures that responded significantly to *num buildings*, response was more complex, with measures worsening initially as the number of buildings increases, then reaching a plateau or even improving slightly as the number of buildings reached maximum. It may be that when the number of buildings was high, the additional location constraints imposed by the physical street pattern on the board limited the number of possible configurations and made the assessment task easier. Alternatively or additionally, since trials with larger neighborhoods were always encountered later in the assessment, participants may simply have been more practiced by the time these larger neighborhoods were encountered.

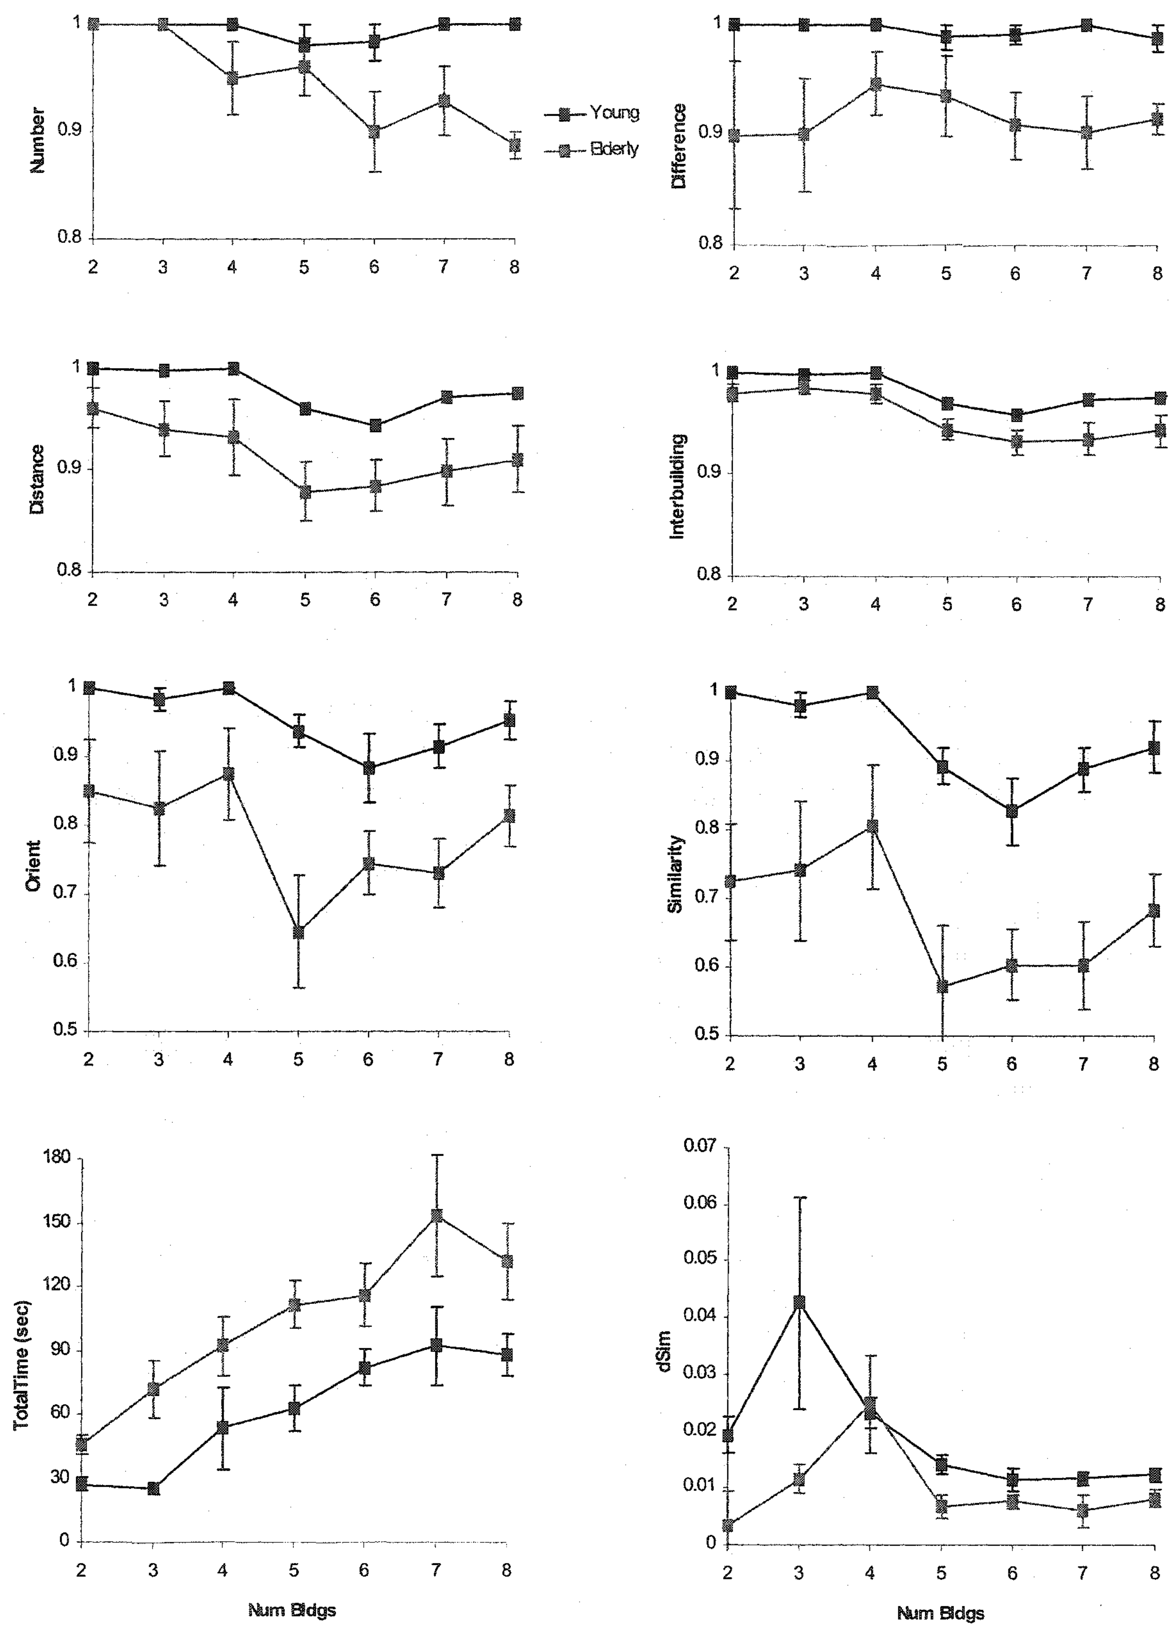

Figure 5.13 – Experimental means and standard errors for the CMP

Only *difference* did not vary significantly as *num buildings* changed. Here the null hypothesis – that the normalized set difference is simply not sensitive to the size of the map participants are attempting to reproduce – likely provides the best explanation of this result.

The effects of *age* and *num buildings* interacted only in the *number*

Table 5.1 – Results of two way ANOVAs in sensitivity study.

Independent Measures	**Dependent Measures**	**ANOVA**
age	*totaltime*	F(1,18)=9.242, p=.007
age	*number*	F(1,18)=14.797, p=.001
age	*difference*	F(1,18)=14.928, p=.001
age	*orientation*	F(1,18)=15.73, p=.001
age	*distance*	F(1,18)=7.2, p=.015
age	*interbuilding*	F(1,18)=10.29, p=.005
age	*similarity*	F(1,18)=18.68, p<.0005
age	*dSim*	F(1,18)=6.759, p=0.018
# bldgs	*totaltime*	F(6,108)=15.432, p<.0005
# bldgs	*number*	F(6,108)=3.400, p=.004
# bldgs	*orientation*	F(6,108)=3.537, p=0.003
# bldgs	*distance*	F(6,108)=6.64, p<.0005
# bldgs	*interbuilding*	F(6,108)=15.789, p<.0005
# bldgs	*similarity*	F(6,108)=5.33, p<.0005
# bldgs	*dSim*	F(6,108)=3.374, p=0.004
age x # bldgs	*number*	F(6,108)=2.884, p=.012

measure. While *num buildings* had little effect on the young, the mapping performance of the elderly dropped significantly by this measure as the number of buildings increased. This is likely due to an age-based difference in recall.

We should note again that because of our need to find the cognitive thresholds of our participants quickly, we ordered experimental trials so that the *num buildings* factor increased steadily. Because of this pointed lack of counterbalancing or randomization in *num buildings,* practice effects might be confounded with the observed effects of *num buildings.*

5.7 Discussion

5.7.1 Confirmations

Our experimentation confirms that the CMP is sensitive to factors known to affect cognitive mapping performance. As expected, the bulk of our results indicate that the elderly perform worse at cognitive mapping than the young. Increasing the size of the map being reproduced can also worsen mapping performance.

We were also pleased with the match of the CMP interface to the mapping task, and its accessibility to the elderly population. Almost all of our participants were able to complete all the trials – and most reported they had fun doing so. This was true whether participants were university students or World War II veterans. Many of our elderly participants also stated, as part of their general impression rather than as an answer to an explicit question, that the CMP tasks were an excellent training exercise for their memory, and that they would like to repeat the test from time to time, as a "brain workout".

We were also gratified to see that our single AD participant was among the worst performers, tentatively indicating possible use of the CMP for palliative care of persons with AD. Much more research is required before this application is realized.

5.7.2 Surprises

We expected that assessment performance would worsen as *num buildings* increased. Instead, *num buildings* had a much more complex impact. While confounding practice effects certainly had an influence on this result, the initial *decrease* in mapping performance as the number of buildings increases (the opposite of a practice effect) leads us to believe that the constraints provided by our tangible street pattern played a larger role. This suggests that mapping difficulty might be controlled in future experiments by varying the proportion of the map used for street cues.

We did not expect the *age* x *num buildings* interaction we saw in our results. It would be interesting to see if performance in the *number* measure also declines for the young as the number of buildings increases further.

5.7.3 What's Next?

While our results indicate great promise for technologies like the CMP, there is much work that remains if its assessment paradigm is to become common in clinical and research settings. First, the measurement sensitivity and reliability of CMP-like tools must be probed further and extensively, with comparisons made to existing assessment techniques and typical score distributions found so that unusual assessment results might quickly be recognized (see Section 6.3). Second, tangible and tabletop based interaction must become cheaper and more reliable, so that newer versions of the CMP will be more cost effective. Cost-effectiveness should also be carefully demonstrated with comparison to a tentative, non-TUI, automatic tool (see Section 6.3 for our work on a reference WIMP-based cognitive mapping assessment tool).

5.8 Conclusion

In this chapter we have presented the CMP, a tangible user interface for the assessment of cognitive mapping ability. In experimentation, the CMP proved to be sensitive to factors known to affect cognitive mapping ability.

We see our work on the CMP as a several-fold accomplishment: fulfilling our expectations in demonstrating the usefulness of TUIs, closely coupled with an applied research problem from a non-computing-oriented domain seeking new ideas and solutions. We see the CMP, like Cognitive Cubes, as a manifestation of our TUI design heuristics: intuitive spatial mapping between interface and task, I/O unification and support for trial-and-error actions. Our heuristics guided our choice when selecting the spatial application, and later throughout the implementation process, seeking and finding direct and intuitive spatial mapping between the CMP's physical interface and the assessment task. We believe that the successful implementation of the CMP reflects extremely well on the relevance of our TUI heuristics to the HCI domain.

We believe that the CMP, like Cognitive Cubes, is an example of the way future innovation in tangible UIs should be closely tied to target applications. Such close relationships will allow researchers to isolate those components of the tangible interface that are strong and those that are weak, in concrete application terms. This approach might also lead to faster adoption of TUIs by the industry, pushing further the current state-of-the-art.

There are many interesting opportunities for improving the CMP's sensitivity. For example the CMP could be used iteratively, with visual feedback given to the participant about the accuracy of the currently reproduced map, enabling the participant to attempt to correct their map. Active or exocentric viewing modes might be explored. The detailed histories of map building compiled by the CMP might be analyzed to find the decision trees formed by participants. Ultimately, the CMP might also prove useful for training and therapeutic applications (see Section 6.3).

Chapter VI
Conclusion and Future Work

6.1 Spatial Tangible User Interfaces for Cognitive Assessment

Tangible user interfaces are a new HCI research domain. The benefits of exploring a new domain are that nearly any path you take is original. The unavoidable snag is that not necessarily every path leads to a meaningful destination.

While greatly benefiting from being among the very first to explore this new domain, we attempted to stay focused on these meaningful destinations. We devoted considerable resources to ethnographic exploration of possible TUI applications. We chose the potential impact on the applied domain as a crucial criterion for the quality of our work. Once we centered our attention on cognitive assessment, all our efforts were shifted to careful, often tedious design, and later experimental demonstration, of useful assessments.

This chapter briefly overviews our work, but more importantly, reemphasizes what we believe is the key, hopefully lasting, two-fold outcome of our work. Firstly our approach to TUI design (manifested in our set of design heuristics) is unique in its simplicity. If we got it right, our approach will have an impact on future TUI design and on related HCI research. Secondly, our demonstrated coupling of TUIs with assessment applications will hopefully result in further efforts, pushing these or similar paradigms towards clinically tested, valuable tools.

6.2 Contributions Revisited

We followed a very simple objective (Section 2.1):

> Show that tangible user interfaces can be *feasibly used in meaningful computer applications.*

We believe we attained this objective—our TUIs effectively tackled meaningful applications, and were among the very first to do so. We designed and implemented innovative TUIs, aimed at assessment of different cognitive abilities. Our TUIs successfully addressed types of interactions that are particularly hard to carry out with the standard UI (for example, 3D construction). Our TUIs were also successful in involving users that are often blocked by the standard UI from these types of interaction (for example, elderly participants).

With our Cognitive Cubes and the CMP we were among the very first to pursue experimental testing of meaningful TUI applications. In the experiments both of our TUIs demonstrated effectiveness as assessment tools, and sensitivity to the cognitive ability they attempted to measure. Furthermore, in Cognitive Cubes we presented the very first experimental testing of a 3D spatial TUI application.

Cognitive Cubes and the CMP emerge from our novel approach to TUIs; an approach manifested in a new taxonomy and new design heuristics, emphasizing the crucial importance we see in intuitive spatial mapping between TUI and task. We employed our heuristics in a thorough overview of TUIs state of the art, accompanied by a new TUI taxonomy that reflects our view of the domain.

In addition, both Cognitive Cubes and the CMP are significant, first-of-their-kind tools in their applied domain of cognitive assessment. Cognitive Cubes and the CMP are the first TUIs to address this domain and the first to be successfully tested in comprehensive experiments. In particular, Cognitive Cubes is the first automatic system for assessment of 3D constructional ability, and the CMP is the first detailed automatic system for assessment of cognitive mapping ability.

6.3 Future Work

From an HCI point of view, we (like others) would like to be able to identify experimentally those situations in which TUIs are better for a task than the standard UI. Cognitive Cubes are an elaborate 3D spatial interface that will be hard, if not impossible, to mimic with the standard UI. This makes its advantage over the standard UI quite obvious, and hence much less of a challenge to demonstrate. On the other hand, the CMP raises the question: "couldn't this be done more simply with a WIMP interface?". We strongly believe that our TUI approach adds considerable ecological validity to the CMP, making it possible to assess diverse populations, such as the elderly and AD patients that might find it hard or even impossible to perform in a WIMP-based assessment. We would like to see this hypothesis experimentally tested.

We are planning a detailed test-comparison experiment of the CMP. This experiment will follow exactly the same assessment process described in Chapter 5, with the exception that the assessment tool will not be a TUI. Instead of the CMP we will use Mapper—a WIMP based

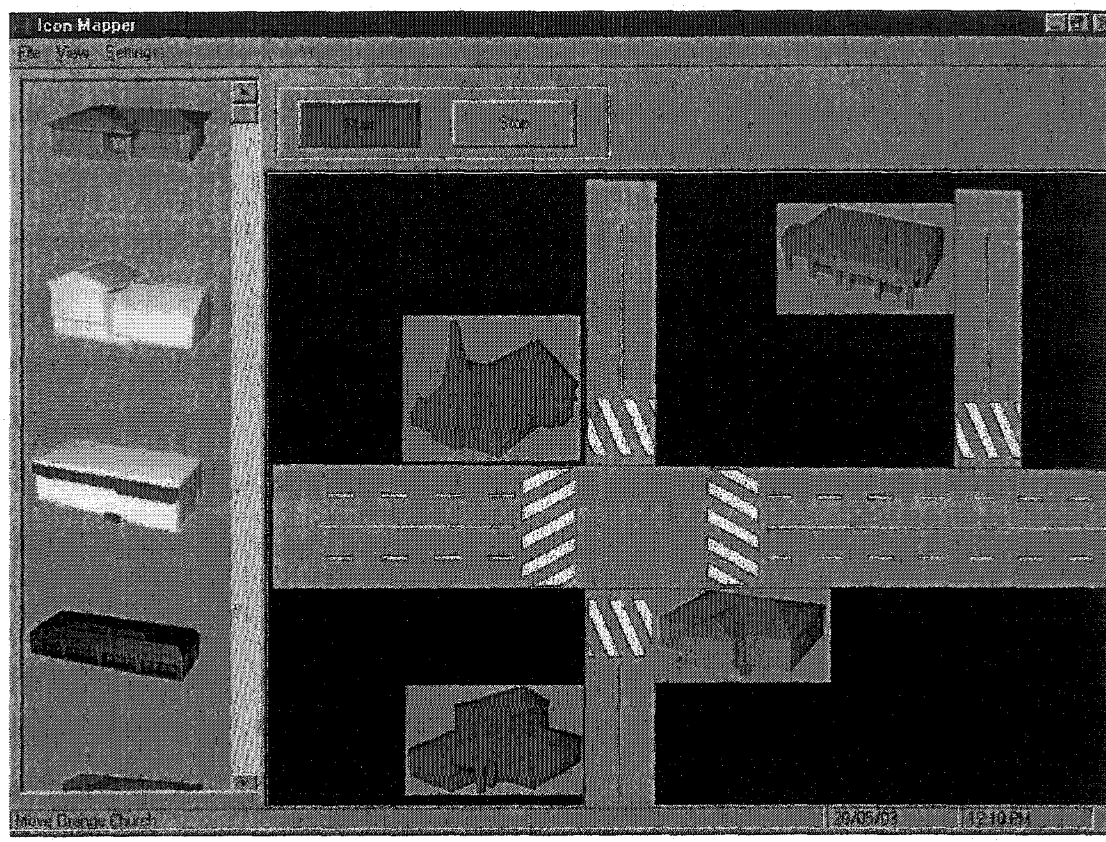

Figure 6.1 – Mapper

cognitive mapping tool that we designed[8] (see Figure 6.1). Mapper enables the participant to input her cognitive map using the standard UI, with all the benefits of automatic assessment. Mapper's results will then be compared to the CMP's results. We hope these test comparison results will more clearly demonstrate the benefits of TUIs over the standard UI, especially in assessment of elderly participants.

Ideally, a comprehensive HCI testing of the effectiveness of our TUI design heuristics would include a comparison to a TUI that was designed

---

[8] Mapper was designed by us and implemented by Angelo Gonzalez from Northwestern University, Illinois.

differently, for example with weaker spatial mapping between the TUI and the task. Future efforts could address a comparison between our systems and such impaired reference systems.

As we mentioned earlier (Chapters 4 and 5), we believe much more analysis could be done with the data collected by the CMP and Cognitive Cubes. Automatic assessment imposes the challenge of dealing with an extremely high density of measurement—a vast amount of data that is naturally sampled by the assessment tool (see Chapter 3). Careful inspection of this data should reveal traces of detailed human behaviors and cognition, namely participants' tactics and strategies when performing the tasks. We believe these behaviors might be revealed by building decision trees tracing a participant's step-by-step line of action. Different strategies might be revealed by aggregating similar, and dissimilar decision trees in clusters, and searching for possible patterns coupling different clusters with different cognitive abilities. We envision such decision tree based measures eventually becoming a useful assessment parameter, much like the traditional time-to-completion.

In order for cognitive assessment ideas like these to thrive, clinical recognition is crucial and something we would very much like to realize. We hope to collaborate with relevant clinicians who might find direct interest in, and benefit from, such cognitive assessment tests. We would like to see inclusions of a significant number of participants with mild AD in these future studies. Our initial plans were to include a balanced number of mild AD participants in our experiments (that is, balanced to the number of healthy elderly and healthy non-elderly). In spite of considerable efforts we could not recruit a sufficient number of participants with mild AD. We believe a future approach to these recruitment difficulties would be the integration of a practicing clinician on our research team, with direct access to AD patients and their caregivers. Clearly, our results would have to be confirmed with a larger number of participants before reaching clinical approval.

We see strong parallels between the cognitive assessment tasks we pursued and potential rehabilitation and training applications. Other than our belief, we have found several direct indications for such a correspondence. For example, in Cognitive Cubes, MRT results were noticeably improved after the Cognitive Cubes session (see Section 4.8.3). Also, participant comments (especially by elderly CMP participants) indicated an improvement of spatial ability and memory.

We can also envision a cognitive mapping training tool based on an enhanced version of the CMP with a larger urban vocabulary (that is, more extensive variety of miniature building models that can be placed on the TUI). Such a tool could be used to train participants in wayfinding in familiar or unfamiliar urban settings. The participant could interact with both the physical and the virtual representations of the environment, and would be expected to construct a correct physical replica of the environment using the TUI. The tool would automatically support the building process with visual feedback highlighting errors and providing hints (for example, showing an error by placing a shadow of a misplaced building in its correct location). Such a trainer could also include navigation components by allowing the participant to explore the environment using a tracked physical avatar (see [139]). As she moves the avatar on the TUI, the participant is presented with virtual, dynamic visual and audio feedback to assist the learning process.

We also believe spatial TUI-based training and rehabilitation tasks could dramatically benefit the visually impaired community. Obviously, a computer interface that is largely non-visual and tactile should hold promise for the visually impaired. Young visually impaired children could benefit from an automatic trainer based for example on ActiveCube (see Section 2.11) when learning basic spatial relations like "on top" or "in the corner". These concepts are not trivial for the visually impaired young to grasp, especially for those completely blind at birth [130,139]. Such a trainer could tie together the participant's physical actions using the TUI with detailed audio feedback interactively summarizing the current spatial relationship.

We designed a Tangible Trainer for Orientation and Mobility (TTOM) in detail, but did not implement it [139]. The TTOM is a TUI based orientation and mobility trainer for the visually impaired (essentially "orientation and mobility" is a different term for wayfinding, see Section 3.2). The TTOM is based on a CMP-like interface, but instead of miniature building models it uses a tactile vocabulary based on the way the environment is perceived by a blind person walking through it with a cane. The TTOM is designed to allow detailed, autonomous learning of a new setting and self-assessment of the resulting cognitive map.

The TTOM objects represent pavement, sidewalks, curbs, ramps, walls and poles which are all common landmarks for a cane-walker [130]. We thought of basing the tangible vocabulary on our own legend, or adopting an existing legend such as the tactile maps used for blind orienteering (see for example [100,102]). The TTOM can be used as a straightforward static tactile map. It can also be used as an interactive tactile map with audio feedback generated according to the position of a tracked avatar on the TUI. Lastly, it can be used as an editable tactile map which interactively supports (by audio feedback) the participants attempts to input a spatial environment from memory.

An interesting possibility might be to use the TTOM as an online tool enabling blind users to independently learn new environments using a computer, much like a seeing person downloads a map off the web [139]. The online text information would include a set of coordinates and object identities, based on the TTOM's vocabulary. Following this textual information the user could construct a detailed physical replica of the online information on her remote TTOM. At the same time the TTOM software would download the digital representation of the environment and supporting audio samples. After that, the manual construction process would be supported by audio feedback from the TTOM, correcting wrongly positioned, or missing objects.

Once the construction has been completed the user could use the TTOM for learning, following the interaction techniques mentioned earlier (that is, the TTOM being used as static, interactive or editable tactile map). Given sufficient miniaturization the TTOM could be a portable, autonomous TUI. It follows that the TTOM could allow a visually impaired user to download a new environment, learn it, and then carry its replica with him when visiting the new environment for the first time (much the same as carrying a map) [139].

Although predictions are often inaccurate, we are obliged to make a few educated guesses about the long-term future of TUIs. We believe TUIs will benefit from several technological trends (see also Section 2.14).

The most obvious of these trends is miniaturization. Once enough attention (and hopefully a "killer-app") is given to TUIs, they should become much smaller and much more detailed. Miniaturization of TUIs should have immediate impact on their expressiveness. Once TUIs reach the size of common Lego blocks, a vast number of design applications will open up, enabling direct, intuitive spatial mapping to numerous tasks. We can envision TUI based prototyping and design applications ranging all the way from home design [139], to mechanical and optical design, and ultimately any scientific and engineering domain that employs physical entities as a means of mediating concepts and models (for example, chemistry, physics, biology, astronomy, etc.).

Further down the road we foresee technology potentially revolutionizing I/O unification in TUIs (see Section 2.2.3). One exciting technological trend is the blurring of physical and virtual with the maturation of mixed reality technology (Section 2.13 [13,116]). We believe TUIs could leap forward when mixed reality technology becomes more common, free of cables and perhaps goggles. With the support of such mixed reality technology TUIs could reach a remarkable visual I/O unification with physical interfaces taking any visual shape and form while in the user's hand.

Once nanotechnology becomes of age it could have a dramatic impact on TUIs and on the level of I/O unification they afford. A TUI based on nano- (or even micro-, or milli-) scale dynamic components could reshape itself at will, enabling extreme physical I/O unification, with digital shapes interactively output into the physical interface.

Future TUI success should bring its extinction as a field. In the long run we hope to see a successful fusion of TUIs into future interaction paradigms, meaning that TUIs would be an integral part of the future notion of "a computer" or even better, simply "a tool". The doubtful reader is invited to read Engelbart's original mouse paper [37] and reflect on times not long ago when a mouse was merely a small mammal with a long tail.

Bibliography

1. R. Aish and P. Noakes, "Architecture Without Numbers - CAAD Based on a 3D Modeling System," *Computer-Aided Design*, vol. 16, no. 6, November, pp. 321-328, 1984.
2. R. Aish, "3D Input for CAAD Systems," *Computer-Aided Design*, vol. 11, no. 2, March, pp. 66-70, 1979.
3. K. Alexander and E. Strommen, "Evolution of the Talking Dinosaur: The (Not So) Natural History of a New Interface for Children," in *Human Factors in Computing Systems CHI 1998.* Los Angeles, CA, USA, April, 1998, pp. 7-8.
4. G. L. Allen, K. C. Kirasic, and R. L. Beard, "Children's Expressions of Spatial Knowledge," *Journal of Experimental Child Psychology*, vol. 48, pp. 114-130, 1989.
5. G. Anagnostou, D. Dewey, and A. T. Patera, "Geometry Defining Processors for Engineering Design and Analysis," *The Visual Computer*, no. 5, pp. 304-315, 1989.
6. D. Anderson, J. L. Frankel, J. W. Marks, A. Agarwala, P. A. Beardsley, J. K. Hodgins, D. L. Leigh, K. Ryall, E. Sullivan, and J. Yedidia, "Tangible Interaction and Graphical Interpretation: A New Approach to Intuitive 3D Modeling," in *Proceedings of SIGGRAPH 2000*, ACM, Ed., 2000, pp. 393-402.
7. D. Anderson, J. Frankel, J. Marks, D. Leigh, K. Ryall, E. Sullivan, and J. Yedidia, "Building Virtual Structures With Physical Blocks," Demo Summary, in *Proceedings of UIST 1999.* Asheville, North Carolina, USA, November, 1999, pp. 71-72.
8. P. Arthur and R. Passini, *Wayfinding: People, Signs, and Architecture.* Toronto: McGraw-Hill Ryerson, 1992.

9. G. Axia, E. M. Peron, and M. R. Baroni, "Environmental Assessment Across the Life Span," in *Environment, Cognition and Action: An Integrated Approach*, T. Garling and G. W. Evans, Eds. New York: Oxford University press, 1991.

10. J. C. Baird, "Studies of the Cognitive Representation of Spatial Relations," *Journal of Experimental Psychology: General*, vol. 108, no. 1, pp. 90-91, 1979.

11. R. Balakrishnan, W. G. Fitzmaurice, G. Kurtenbach, and K. Singh, "Exploring Interactive Curve and Surface Manipulation Using a Bend and Twist Sensitive Input Strip," in *Proceedings of the ACM Symposium on Interactive 3D Graphics (I3DG'99)*. New York: ACM, 1999, pp. 111-118.

12. M. Beaudouin-Lafon, "Instrumental Interaction: An Interaction Model for Designing Post-WIMP User Interfaces," http://delivery.acm.org/10.1145/340000/332473/p446-beaudouin-lafon.pdf?key1=332473&key2=9934254401&coll=portal&dl=ACM&CFID=7517975&CFTOKEN=94592795, in *Human Factors in Computing Systems CHI 2000*. ACM, 2000, pp. 446-453.

13. M. Billinghurst, H. Kato, and I. Poupyrev, "The MagicBook - Moving Seamlessly Between Reality and Virtuality," *IEEE Computer Graphics and Applications*, May/June, pp. 6-8, 2001.

14. E. Boring, "Intelligence as the Tests Test It," in *The New Republic*, June, 1923, pp. 35-37.

15. D. A. Bowman and E. T. Davis, "Maintaining Spatial Orientation During Travel in an Immersive Virtual Environment," *Presence: Teleoperators and Virtual Environments*, vol. 8, no. 6, December, pp. 618-632, 1999.

16. S. Brave, H. Ishii, and A. Dahley, "Tangible Interfaces for Remote Collaboration and Communication," in *Proceedings of CSCW '98*, 1998, pp. 1-10.

17. J. Brockman, http://www.edge.org/digerati/simonyi/simonyi_p1.html. Edge Foundation, Inc.

18. F. W. Bruns and V. Brauer, "Bridging the Gap Between Real and Virtual Modeling - A New Approach to Human-Computer Interaction," http://www.artec.uni-bremen.de/field1/rugams/papers/texas/texas.html, in *Proceedings of the 2nd IFIP 5.10 Workshop on Virtual Prototyping.* Arlington Texas, May, 1996.

19. W. E. Carlson, "CGI Historical Timeline," http://www.accad.ohio-state.edu/~waynec/history/timeline.html, January 20, 2003.

20. A. Chang, B. Resner, B. Koerner, X. Wang, and H. Ishii, "LumiTouch: An Emotional Communication Device," in *Human Factors in Computing Systems CHI 2001*, ACM, Ed. Seattle, Washington, USA, 2001, pp. 313-314.

21. A. Cho, R. Lederer, E. Sharlin, B. A. Watson, L. Liu, and S. Sutphen, "Designing Virtual Worlds - the Cognitive Map Probe," http://www.cs.nwu.edu/~watsonb/school/docs/ski02.cmp.pdf, in *Proc. 13th Western Canadian Computer Graphics Symposium.* Silver Star, BC, Canada, March, 2002.

22. G. D. Cohen, N. Ring, and M. de Sousa, "Creativity: As We Grow Old," http://www.ablestable.com/resources/library/articles/clife/clife012.htm. AbleStable, 2002-2003.

23. J. Cohen, M. Withgott, and P. Piernot, "Logjam: A Tangible Multi-Person Interface for Video Logging," in *Human Factors in Computing Systems CHI 1999.* Pittsburgh, USA: ACM, 1999, pp. 128-135.

24. H. A. Colle and G. B. Reid, "The Room Effect: Metric Spatial Knowledge of Local and Separated Regions," *Presence: Teleoperators and Virtual Environments*, vol. 7, no. 2, April, pp. 116-129, 1998.
25. E. Cornell, Private Communication. University of Alberta, 2000-2002.
26. G. Crampton Smith, "The Hand That Rocks the Cradle," *ID Magazine*, May/June, pp. 60-65, 1995.
27. R. P. Darken and T. Allard, "Spatial Orientation and Wayfinding in Large-Scale Virtual Spaces II," *Presence: Teleoperators and Virtual Environments*, vol. 8, no. 6, December, pp. iii-vii, 1999.
28. R. P. Darken, T. Allard, and L. B. Achille, "Spatial Orientation and Wayfinding in Large-Scale Virtual Spaces: An Introduction," *Presence: Teleoperators and Virtual Environments*, vol. 7, no. 2, April, pp. 101-107, 1998.
29. R. P. Darken and W. P. Banker, "Navigating in Natural Environments: A Virtual Environment Training Transfer Study," http://www-npsnet.cs.nps.navy.mil/darken/Publications/naturalenvironments.pdf, in *Proceedings of VRAIS '98*, 1998, pp. 12-19.
30. R. P. Darken and S. R. Goerger, "The Transfer of Strategies from Virtual to Real Environments: An Explanation for Performance Differences?" http://www-npsnet.cs.nps.navy.mil/darken/Publications/StrategyTransfer.pdf, in *Virtual Worlds and Simulation (VWsim'99)*, 2. San Francisco, CA, USA, January, 1999, Vol. 31, pp. 159-164.
31. R. P. Darken and B. Peterson, "Spatial Orientation, Wayfinding, and Representation," in *Handbook of Virtual Environments: Design, Implementation, and Applications (Human Factors and Ergonomics)*, K. Stanney, Ed. Lawrence Erlbaum Associates, 2002.

32. H. B. Deets, "Aging and Technology: The Convergence of Two Revolutions," *CyberPsychology and Behavior*, vol. 2, no. 6, December, pp. 501-503, 1999.

33. A. J. Dix, J. E. Finlay, G. D. Abowd, and Beale R, *Human-Computer Interaction*, 2nd. Prentice Hall, 1998.

34. S. Don and J. Duncan, *The Making of Jurassic Park.* New York: Ballantine Books, 1993.

35. P. Dourish, *Where The Action Is.* Cambridge, Massachusetts: The MIT Press, 2001.

36. C. Emlet, J. L. Crabtree, V. Condon, and L. Treml, *In-Home Assessment of Older Adults - An Interdisciplinary Approach.* Gaitherburg, Md: Aspen Publishers, 1996.

37. W. English, D. C. Engelbart, and M. Berman, "Display-Selection Techniques for Text Manipulation," http://www.bootstrap.org/augment/AUGMENT/133184.html, *IEEE Transactions on Human Factors in Electronics*, vol. HFE-8, no. 1, March, pp. 5-15, 1967.

38. C. Esposito and W. B. Paley, "Of Mice and Monkeys: A Specialized Input Device for Virtual Body Animation," in *Symposium on Interactive 3D Graphics.* Monterey, CA, USA, 1995, pp. 109-114.

39. W. G. Fitzmaurice, H. Ishii, and W. Buxton, "Bricks: Laying the Foundations for Graspable User Interfaces," in *Human Factors in Computing Systems CHI 1995.* ACM, May, 1995, pp. 442-449.

40. W. G. Fitzmaurice, "Graspable User Interfaces," Ph.D. Dissertation, http://www.dgp.toronto.edu/people/GeorgeFitzmaurice/thesis/Thesis.gf.html. University of Toronto, 1996.

41. W. G. Fitzmaurice, R. Balakrishnan, and G. Kurtenbach, "Sampling, Synthesis, and Input Devices," *Communications of the ACM*, vol. 42, no. 8, August, pp. 54-63, 1999.

42. M. Fjeld, M. Bichsel, and M. Rauterberg, "BUILD-IT: An Intuitive Design Tool Based on Direct Object Manipulation," http://www.iha.bepr.ethz.ch/pages/leute/fjeld/publications/GW.pdf, in *Gesture and Sign Language in Human Computer Interaction, Lecture Notes in Artificial Intelligence*, I. Wachsmut and M. Frolich, Eds., 1998, Vol. 1371, pp. 297-308.

43. M. Fjeld, M. Bichsel, and M. Rauterberg, "BUILD-IT: A Brick-Based Tool for Direct Interaction," http://www.iha.bepr.ethz.ch/pages/leute/fjeld/publications/EPCEbuildit.pdf, in *Engineering Psychology and Cognitive Ergonomics (EPCE), Vol. 4*, D. Harris, Ed. Hampshire: Ashgate, 1999, pp. 205-212.

44. M. Fjeld, N. Ironmonger, S. G. Schär, and H. Krueger, "Design and Evaluation of Four AR Navigation Tools Using Scene and Viewpoint Handling," in *Proc. Conference of Human-Computer Interaction, INTERACT2001*. Tokyo, Japan, 2001, pp. 214-223.

45. M. Fjeld, N. Ironmonger, F. Voorhorst, M. Bichsel, and M. Rauterberg, "Camera Control in a Planar, Graspable Interface," http://www.iha.bepr.ethz.ch/pages/leute/fjeld/publications/iaested.pdf, in *17th IASTED International Conference Applied Informatics (AI'99)*, 1999, pp. 242-245.

46. M. Fjeld, K. Lauche, M. Bichsel, F. Voorhorst, H. Krueger, and M. Rauterberg, "Physical and Virtual Tools: Activity Theory Applied to the Design of Groupware," in *Computer Supported Cooperative Work*. Netherlands: Kluwer Academic Publishers, 2002, pp. 153-180.

47. M. Fjeld, K. Lauche, S. Dierssen, M. Bichsel, and M. Rauterberg, "BUILD-IT: A Brick-Based, Integral Solution Supporting Multidisciplinary Design Tasks," in *Designing Effective and Usable Multimedia Systems*, http://www.iha.bepr.ethz.ch/pages/leute/fjeld/publications/IFIP1

3.2.pdf, A. Sutcliffe, J. Ziegler, and P. Johnson, Eds. Boston: Kluwer Academic, 1998, pp. 122-133.

48. M. Fjeld, S. G. Schär, D. Signorello, and H. Krueger, "Alternative Tools for Tangible Interaction: A Usability Evaluation," in *Proceedings of the IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR 2002)*, 2002, pp. 157-166.

49. M. F. Folstein, S. E. Folstein, and P. R. McHugh, "Mini-Mental State: A Practical Method for Grading the State of Patients for the Clinician," *Journal of Psychiatric Research*, vol. 12, pp. 189-198, 1975.

50. A. Francetic and H. Shwe, "Smarter Play for Smarter Toys: The Benefits of Technology-Enhanced Play," http://www.gdconf.com/archives/2000/shwe.doc. _2000 Game Design Proceedings, 2000.

51. J. H. Frazer, "Three-Dimensional Data Input Devices," in *Computers/Graphics in the Building Process*. Washington: National Academy of Sciences, 1982.

52. J. H. Frazer, "Use of Simplified Three - Dimensional Computer Input Devices to Encourage Public Participation in Design," in *Computer Aided Design 82, Conference Proceedings*. Butterworth Scientific, 1982, pp. 143-151.

53. J. H. Frazer, *An Evolutionary Architecture*. Architectural Association, 1995.

54. J. H. Frazer, J. M. Frazer, and P. A. Frazer, "Intelligent Physical Three-Dimensional Modelling System," in *Computer Graphics '80, Conference Proceedings*. Online Publications, 1980, pp. 359-370.

55. J. H. Frazer, J. M. Frazer, and P. A. Frazer, "New Developments in Intelligent Modelling," in *Computer Graphics '81, Conference Proceedings*. Online Publications, 1981, pp. 139-154.

56. J. J. Gibson, *The Ecological Approach to Visual Perception*. London: Erlenbaum, L, 1986.

57. S. R. Goerger, R. P. Darken, M. Boyd, T. Gagnon, S. Liles, J. Sullivan, and J. Lawson, "Spatial Knowledge Acquisition from Maps and Virtual Environments in Complex Architectural Spaces," http://www-npsnet.cs.nps.navy.mil/darken/Publications/HerrmannHall.pdf. Proceedings of the 16th Applied Behavioral Sciences Symposium. U.S. Air Force Academy, Colorado Springs, CO., USA, April, 1998, pp. 6-10.

58. R. G. Golledge, "Cognition of Physical and Built Environments," in *Environment, Cognition and Action: An Integrated Approach*, T. Garling and G. W. Evans, Eds. New York: Oxford UP, 1991.

59. G. M. Gorbet, "Beyond Input Devices, A New Conceptual Framework for the Design of Physical-Digital Objects," Masters thesis, http://tangible.media.mit.edu:2727/thesis/thesis.pdf. MIT, 1998.

60. G. M. Gorbet, M. Orth, and H. Ishii, "Triangles: Tangible Interface for Manipulation and Exploration of Digital Information Topography," in *Human Factors in Computing Systems CHI 1998*, April, 1998, pp. 49-56.

61. M. Gorbet and M. Orth, "Triangles: A Physical/Digital Construction Kit," in *Proceedings of Designing Interactive Systems: Processes, Practices Methods and Techniques (ACM DIS '97)*, ACM, Ed. Amsterdam, August, 1997, pp. 125-128.

62. T. Grossman, R. Balakrishnan, and K. Singh, "An Interface for Creating and Manipulating Curves Using a High Degree-Of-Freedom Input Device," http://www.dgp.toronto.edu/~ravin/videos/chi2003_shapetape.mpg, in *Human Factors in Computing Systems CHI 2003*. ACM, 2003.

63. G. Groth-Marnat, *Handbook of Psychological Assessment*, 3rd. New York: John Wiley & Sons, 1997.

64. R. A. Heinlein, "Waldo," in *Waldo and Magic, Inc.* NY: Garden City: Doubleday, 1950.

65. J. F. Herman, "Children's Cognitive Maps of Large-Scale Spaces: Effects of Exploration, Direction, and Repeated Experience," *Journal of Experimental Child Psychology*, vol. 29, pp. 126-143, 1980.

66. C. Hertzog and B. Rypma, "Age Differences in Components of Mental Rotation Task Performance," *Bulletin of the Psychonomic Society*, vol. 29, no. 3, pp. 209-212, 1991.

67. K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell, "Passive Real-World Interface Props for Neurosurgical Visualization," http://research.microsoft.com/users/kenh/papers/Props.pdf, in *Human Factors in Computing Systems CHI 1994*. ACM, 1994, pp. 452-458.

68. J. H. J. Howard and S. M. Kerst, "Memory and Perception of Cartographic Information for Familiar and Unfamiliar Environments," *Human Factors*, vol. 23, no. 4, pp. 495-504, 1981.

69. H. Ishii and B. Ullmer, "Tangible Bits: Towards Seamless Interfaces Between People, Bits and Atoms," in *Human Factors in Computing Systems CHI 1997*, March, 1997, pp. 234-241.

70. Y. Itoh, Y. Kitamura, M. Kawai, and F. Kishino, "Implementation of ActiveCube for Multi-Modal Interaction," in *Proc. Conference of Human-Computer Interaction, INTERACT2001*. Tokyo, Japan, 2001, pp. 682-683.

71. R. J. K. Jacob, H. Ishii, G. Pangaro, and J. Patten, "A Tangible Interface for Organizing Information Using a Grid," http://tangible.media.mit.edu/papers/Senseboard_CHI02/Senseboard_CHI02.pdf, in *Human Factors in Computing Systems CHI 2002*. Minneapolis, USA: ACM, 2002.

72. M. P. Johnson, A. Wilson, C. Kline, B. Blumberg, and A. Bobick, "Sympathetic Interfaces: Using a Plush Toy to Direct Synthetic Characters," http://web.media.mit.edu/~bruce/Site01.data/SympatheticInterfaces.pdf, in *Human Factors in Computing Systems CHI 1999.* Pittsburgh, May, 1999.
73. H. Kato, M. Billinghurst, I. Poupyrev, K. Imamoto, and K. Tachibana, "Virtual Object Manipulation on a Table-Top AR Environment," in *Proceedings of ISAR 2000*, 2000.
74. D. G. Kessler, D. A. Bowman, and L. F. Hodges, "The Simple Virtual Environment Library, an Extensible Framework for Building VE Applications," *Presence: Teleoperators and Virtual Environments*, vol. 9, no. 2, April, pp. 187-208, 2000.
75. M. J. Kincade, "Visuospatial Perception Disturbances in Alzheimer's Disease," Masters thesis. University of Alberta, 1996.
76. K. C. Kirasic, "Age Sensitive Spatial Abilities, Environmental Knowledge, and Wayfinding: Mediated Relationships and Missing Links," *Spatial Cognition and Computation*, 2001.
77. D. Kirsh and P. Maglio, "On Distinguishing Epistemic from Pragmatic Action," *Cognitive Science*, vol. 18, pp. 513-549, 1994.
78. Y. Kitamura, "ActiveCube," http://www-human.ist.osaka-u.ac.jp/ActiveCube/.
79. Y. Kitamura, Y. Itoh, and F. Kishino, "Real-Time 3D Interaction with ActiveCube," Short Talks, in Extended Abstracts, in *Human Factors in Computing Systems CHI 2001.* Seattle, Washington, USA, March, 2001, pp. 355-356.
80. Y. Kitamura, Y. Itoh, T. Masaki, and F. Kishino, "ActiveCube: A Bi-Directional User Interface Using Cubes," in *KES 2000: Fourth International Conference on Knowledge-Based Intelligent Engineering*

Systems & Allied Technologies. University of Brighton, UK, August, 2000, pp. 99-102.

81. K. Kobayashi, J. Patten, D. Chak, J. Alonso, , and H. Ishii, "IP Network Design Workbench," http://tangible.media.mit.edu/projects/ipnet_workbench/ipnet_workbench.html. MIT, Media Lab, Tangible Media Group, 2002.
82. J. J. Koenderink, *Solid Shape*. Cambridge, MA: The MIT Press, 1990.
83. G. Koh, V. Von, and E. Thomas, "Use of Virtual Environments for Acquiring Configurational Knowledge About Specific Real-World Spaces: Preliminary Experiment," *Presence: Teleoperators and Virtual Environments*, vol. 8, no. 6, December, pp. 632-657, 1999.
84. G. Kurtenbach and R. Balakrishnan, "HCI and the Fall of the Virtual/Physical Wall," Paper draft - unpublished manuscript, 1999.
85. P. Larson, A. A. Rizzo, J. G. Buckwalter, A. van Rooyen, K. Kratz, U. Neumann, C. Kesselman, M. Thiebaux, and C. van der Zaag, "Gender Issues in the Use of Virtual Environments," *CyberPsychology and Behavior*, vol. 2, no. 2, 1999.
86. J. Lee, B. Dunn, S. Ren, V. Su, and H. Ishii, "GeoSCAPE: Designing a Reconstructive Tool for Field Archaeological Excavation," http://tangible.media.mit.edu/papers/GeoSCAPE_CHI01/GeoSCAPE_CHI01.pdf, in *Human Factors in Computing Systems CHI 2001; Extended abstracts*. Seattle, Washington, USA: ACM, 2001.
87. J. Lee, V. Su, S. Ren, and H. Ishii, "HandSCAPE: A Vectorizing Tape Measure for On-Site Measuring Applications," in *Human Factors in Computing Systems CHI 2000*. ACM, 2000, pp. 137-144.
88. Lego, "Music Builder Composer," http://www.lego.com/eng/explore/musicbuilder.asp.

89. M. D. Lezak, *Neuropsychological Assessment.* New York: Oxford UP, 1983.

90. L. Liu, Private Communication. University of Alberta, 1999-2002.

91. L. Liu, L. Gauthier, and S. Gauthier, "Spatial Disorientation in Persons with Early Senile Dementia of the Alzheimer Type," *The American Journal of Occupational Therapy*, vol. 45, no. 1, January, pp. 67-74, 1991.

92. L. Liu, L. Gauthier, and S. Gauthier, "The Functional Spatial Abilities Questionnaire for Use with Persons Who Are in the Early Stages of Alzheimer Disease: Preliminary Data on Reliability and Validit," *Canadian Journal of Occupational Therapy*, vol. 63, pp. 222-233, 1996.

93. L. Liu, B. A. Watson, and M. Miyazaki, "VR for the Elderly: Quantitative and Qualitative Differences in Performance with a Driving Simulator," *CyberPsychology and Behavior*, vol. 2, no. 6, pp. 567-576, 1999.

94. M. Lumbreras and J. Sanchez, "Interactive 3D Sound Hyperstories for Blind Children," Human Factors in Computing Systems CHI 1999, May, 1999, pp. 318-325.

95. W. E. Mackay, A. L. Fayard, L. Frobert, and L. Medini, "Reinventing the Familiar: Exploring an Augmented Reality Design Space for Air Traffic Control," in *Human Factors in Computing Systems CHI 1998.* Los Angeles, CA, USA: ACM, 1998, pp. 558-565.

96. D. R. McGee, P. R. Cohen, R. M. Wesson, and S. Horman, "Comparing Paper and Tangible, Multimodal Tools," in *Human Factors in Computing Systems CHI 2002.* Minneapolis, Minnesota, USA: ACM, 2002, pp. 407-414.

97. T. P. McNamara, "How Are the Locations of Objects in the Environment Represented in Memory?" in *Spatial Cognition III*, http://www.psy.vanderbilt.edu/faculty/mcnamara/lab/visual/SCII

I.pdf, C. Freksa, W. Brauer, C. Habel, and K. F. Wender, Eds. Berlin: Springer, In press.

98. M. Meehan, B. Insko, M. Whitton, and F. P. Brooks, "Physiological Measures of Presence in Stressful Virtual Environments," in *Proceedings of SIGGRAPH 2002*. San Antonio, Texas: ACM, 2002, pp. 645-652.

99. D. Meichenbaum and R. Cameron, "Issues in Cognitive Assessment: An Overview," in *Cognitive Assessment*, T. V. Merluzzi, C. R. Glass, and M. Genest, Eds. New York, USA: The Guilford Press, 1981.

100. I. Morisbak, "Orienteering: What It is and How It Can Be Provided as an Activity for Handicappers," http://www.iub.edu/~outdoor/bponline/bp1982/bp82mori.doc, *The Bradford Papers Online*, 1982. Institute for Outdoor Leadership and Education.

101. T. Murakami, K. Hayashi, K. Oikawa, and N. Nakajima, "DO-IT: Deformable Objects as Input Tools," in *Human Factors in Computing Systems CHI 1995, Conference Companion*. Denver, 1995, pp. 87-88.

102. G. Murphy, Private Communication. Community Services, Edmonton, 2000-2001.

103. Neurosmith, "Music Blocks," http://www.neurosmith.com/pd_mblocks.php.

104. W. Newman and P. Wellner, "A Desk Supporting Computer-Based Interaction with Paper Documents," in *Human Factors in Computing Systems CHI 1992*. ACM, 1992, pp. 587-592.

105. J. Nielsen, "Enhancing the Explanatory Power of Usability Heuristics," http://doi.acm.org/10.1145/191666.191729, in *Human Factors in Computing Systems CHI 1999*. ACM, 1994, pp. 152-158.

106. D. A. Norman, *The Psychology of Everyday Things*. BasicBooks, 1988.
107. D. A. Norman, *The Invisible Computer*. MIT Press, 1998.
108. G. Pangaro, D. Maynes-Aminzade, and H. Ishii, "The Actuated Workbench: Computer-Controlled Actuation in Tabletop Tangible Interfaces," http://tangible.media.mit.edu/papers/Actuated_Workbench_UIST02/uist2002actuatedworkbench.pdf, in *Proceedings of UIST 2002*. ACM, 2002.
109. E. Patrick, D. Cosgrove, A. Slavkovic, J. Rode, T. Verratti, and G. Chiselko, "Using a Large Projection Screen as an Alternative to Head-Mounted Displays for Virtual Environments," in *Human Factors in Computing Systems CHI 2000*, April, 2000, pp. 478-485.
110. J. Patten, H. Ishii, J. Hines, and G. Pangaro, "Sensetable: A Wireless Object Tracking Platform for Tangible User Interfaces," in *Human Factors in Computing Systems CHI 2001*. Seattle, Washington, USA: ACM, 2001.
111. J. Patten, B. Recht, and H. Ishii, "Audiopad: A Tag-Based Interface for Musical Performance," http://tangible.media.mit.edu/papers/audiopad_NIME02/audiopad_NIME02.pdf, in *NIME 02 - New Interfactor Musical Expression*. Dublin, Ireland, 2002.
112. M. Peters, B. Laeng, K. Latham, M. Johnson, R. Zaiyouna, and C. Richardson, "A Redrawn Vandenberg & Kuse Mental Rotation Test: Different Versions and Factors That Affect Performance," *Brain and Cognition*, vol. 28, pp. 39-58, 1995.
113. K. Pezdek and G. W. Evans, "Visual and Verbal Memory for Objects and Their Spatial Locations," *Journal of Experimental Psychology: Human Learning and Memory*, vol. 5, no. 4, pp. 360-373, 1979.

114. J. Piaget and B. Inhelder, *The Child's Conception of Space.* New-York: The Humanities Press INC., 1956, pp. 419-446.

115. B. Piper, C. Ratti, and H. Ishii, "Illuminating Clay: A 3-D Tangible Interface for Landscape Analysis," in *Human Factors in Computing Systems CHI 2002.* Minneapolis, Minnesota, USA: ACM, 2002, pp. 355-362.

116. I. Poupyrev, D. S. Tan, M. Billinghurst, H. Kato, H. Regenbrecht, and N. Tetsutani, "Developing a Generic Augmented Reality Interface," *IEEE Computer*, vol. 35, no. 3, March, pp. 44-50, 2002.

117. R. Rasker, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs, "The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays," in *Proceedings of SIGGRAPH 98*, 1998.

118. M. Rauterberg, T. Mauch, and R. Stebler, "What is a Promising Candidate for the Next Generation of Interface Technology?" http://www.ipo.tue.nl/homepages/mrauterb/publications/DPD96 paper.pdf, in *Proceedings of '5th International Conference INTERFACE to Real & Virtual Worlds,'* 1996, pp. 95-103.

119. B. Reisberg, "Functional Assessment Staging (FAST)," http://www.bigtreemurphy.com/Reisberg%20FAST.htm, *Psychopharmacology Bulletin*, vol. 24, pp. 653-659, 1988.

120. J. Rekimoto and M. Saitoh, "Augmented Surfaces: A Spatially Continuos Work Space for Hybrid Computing Environments," in *Human Factors in Computing Systems CHI 1999.* Pittsburgh, USA: ACM, 1999, pp. 378-385.

121. J. Rekimoto, B. Ullmer, and H. Oba, "DataTiles: A Modular Platform for Mixed Physical and Graphical Interactions," in *Human Factors in Computing Systems CHI 2001.* Seattle, WA, USA: ACM, 2001, pp. 269-276.

122. G. Riva, A. A. Rizzo, D. Alpini, E. A. Attree, E. Barbieri, L. Bertella, J. G. Buckwalter, R. C. Davies, L. Gamberini, G. Johansson, N. Katz, S. Marchi, L. Mendozzi, E. Molinari, L. Pugnetti, D. F. Rose, and P. L. Weiss, "Virtual Environments in The Diagnosis, Prevention and Intervention of Age-Related Diseases: A Review of VR Scenarios Proposed in the EC VETERAN Project," *CyberPsychology and Behavior*, vol. 2, no. 6, pp. 577-593, 1999.

123. A. A. Rizzo, T. Bowerly, J. G. Buckwalter, M. T. Schultheis, Matheis R, C. Shahabi, U. Neumann, D. Kirsh, L. Kim, and M. Sharifzadeh, "Virtual Environments for the Assessment of Attention and Memory Processes: Results from the Virtual Classroom and Office," http://dimlab.usc.edu/DocsDemos/ICDVRAT2002.pdf, in *Proceedings of the International Conference on Disability, Virtual Reality and Associated Technology*. Vesaprem, Hungary, September, 2002.

124. A. A. Rizzo and J. G. Buckwalter, "Virtual Reality and Cognitive Assessment and Rehabilitation: The State of the Art," in *Virtual Reality in Neuro-Psycho-Physiology: Cognitive, Clinical and Methodological Issues in Assessment and Rehabilitation*, G. Riva, Ed. Amsterdam: IOS Press, 1997, pp. 123-146.

125. A. A. Rizzo and J. G. Buckwalter, "Special Issue: Virtual Reality Applications in Neuropsychology. Guest Editor's Introduction," *Presence: Teleoperators and Virtual Environments*, vol. 10, no. 4, August, 2001.

126. A. A. Rizzo, J. G. Buckwalter, J. S. McGee, T. Bowerly, C. van der Zaag, U. Neumann, M. Thiebaux, L. Kim, J. Pair, and C. Chua, "Virtual Environments for Assessing and Rehabilitating Cognitive/Functional Performance - A Review of Projects at the USC Integrated Media Systems Center," *Presence: Teleoperators and Virtual Environments*, vol. 10, no. 4, August, pp. 359-374, 2001.

127. A. A. Rizzo, J. G. Buckwalter, U. Neumann, C. Chua, A. van Rooyen, P. Larson, K. Kratz, C. Kesselman, M. Thiebaux, and L. Humphrey, "Virtual Environments for Targeting Cognitive Processes: An Overview of Projects at the University of Southern California," *CyberPsychology and Behavior*, vol. 2, no. 2, pp. 89-100, 1999.

128. A. A. Rizzo, J. G. Buckwalter, U. Neumann, C. Kesselman, M. Thiebaux, P. Larson, and A. van Rooyen, "The Virtual Reality Mental Rotation/Spatial Skills Project: Preliminary Findings," *CyberPsychology and Behavior*, vol. 1, no. 2, pp. 107-113, 1998.

129. A. A. Rizzo, G. Burdea, L. F. Hodges, B. Wiederhold, M. T. Schultheis, M. Wiederhold, and D. Strickland, "Assessment, Rehabilitation and Therapy Applications Using VR," in *Proceedings of the IEEE Virtual Reality 2002*. Orlando, FL: IEEE, 2002, p. 301.

130. D. Rorabeck, Private Communication. The Canadian National Institute for the Blind (CNIB), Edmonton, 1999-2001.

131. D. F. Rose, E. A. Attree, B. M. Brooks, and T. K. Andrews, "Learning and Memory in Virtual Environments: A Role in Neurorehabilitation? Questions (and Occasional Answers) from the University of East London," *Presence: Teleoperators and Virtual Environments*, vol. 10, no. 4, August, pp. 345-358, 2001.

132. B. Rosenthal, "Functional Vision Changes in Aging," in *Geriatric Rehabilitation Manual*, T. L. Kauffman, Ed. New York: Churchill Livingstone, 1999.

133. R. A. Ruddle, S. J. Payne, and D. M. Jones, "Navigating Large-Scale 'Desk-Top' Virtual Buildings: Effects of Orientation Aids and Familiarity," *Presence: Teleoperators and Virtual Environments*, vol. 7, no. 2, April, pp. 179-193, 1998.

134. J. Schaeffer, Private Communication. University of Alberta, 1999-2002.

135. K. Schafer, V. Brauer, and F. W. Bruns, "A New Approach to Human-Computer Interaction - Synchronous Modeling in Real and Virtual Spaces," http://www.artec.uni-bremen.de/field1/rugams/papers/DIS97Paper/DISPaper.html, in *Proceedings of Designing Interactive Systems: Processes, Practices Methods and Techniques (ACM DIS '97)*. Amsterdam, NL: ACM, 1997.

136. M. T. Schultheis and A. A. Rizzo, "The Application of Virtual Reality Technology for Rehabilitation," *Rehabilitation Psychology*, vol. 46, no. 3, pp. 296-311, 2001.

137. E. Sharlin, Y. Itoh, B. A. Watson, Y. Kitamura, S. Sutphen, and L. Liu, "Cognitive Cubes: A Tangible User Interface for Cognitive Assessment," http://www.cs.nwu.edu/~watsonb/school/docs/chi02.pdf, in *Human Factors in Computing Systems CHI 2002*. Minneapolis, US: ACM, April, 2002, pp. 347-354.

138. E. Sharlin and B. A. Watson, "BLOXELS - a "Natural" 3D Modeler; New Tangible Paradigms to Shape Input and Design," Preliminary Research Proposal. University of Alberta, 1999.

139. E. Sharlin and B. A. Watson, "New Paradigms in Tangible Interfaces," Research Proposal (Candidacy Document). University of Alberta, Edmonton, Alberta, July, 2000.

140. E. Sharlin, B. A. Watson, L. Liu, S. Sutphen, R. Lederer, P. Figueroa, and J. Frazer, "Assessmentville - a Preview," http://www.cs.nwu.edu/~watsonb/school/docs/gi01.abstract.assess.pdf. Graphics Interface 2001. Ottawa, Canada, June, 2001.

141. E. Sharlin, B. A. Watson, L. Liu, S. Sutphen, R. Lederer, P. Figueroa, and J. Frazer, "Printing Virtual Reality Interfaces: The Cognitive Map Probe," http://www.cs.nwu.edu/~watsonb/school/docs/vr01.cmp.pdf, in

The Future of VR and AR Interfaces: multi-modal, humanoid, adaptive and intelligent. Proc. Workshop IEEE Virtual Reality 01.' GMD Report 138, W. Broll and L. Schafer, Eds. GMD - German National Research Center for Information Technology., July, 2001, pp. 40-42.

142. E. Sharlin, B. A. Watson, S. Sutphen, R. Lederer, P. Figueroa, and J. Frazer, "3D Computer Interaction Using Physical Objects: Exploration of Tangible User Interfaces," http://www.cs.nwu.edu/~watsonb/school/docs/leonardo.pdf, *Leonardo Electronic Almanac*, vol. 9, no. 7, August, 2001.
143. A. L. Shelton and T. P. McNamara, "Systems of Spatial Reference in Human Memory," *Cognitive Psychology*, vol. 43, pp. 274-310, 2001.
144. R. N. Shepard and J. Metzler, "Mental Rotation of Three-Dimensional Objects," *Science*, vol. 171, pp. 701-703, 1971.
145. B. Shneiderman, "Direct Manipulation: A Step Beyond Programming Languages," *IEEE Computer*, vol. 16, no. 8, August, pp. 57-69, 1983.
146. A. W. Siegel, J. F. Herman, G. L. Allen, and K. C. Kirasic, "The Development of Cognitive Maps of Large- and Small-Scale Spaces," *Child Development*, vol. 50, pp. 582-585, 1979.
147. Sierras Studios, "The Official Half-Life Web-Site," http://www.sierrastudios.com/games/half-life/.
148. M. J. Sinclair, "The Haptic Lens," http://www.me.psu.edu/me415/SPRING02/Microsoft/index.html, in *Visual Proceedings of SIGGRAPH' 97*, ACM, Ed. Los Angeles: ACM, 1997, p. 179.
149. "Sketchpad: The First Interactive Computer Graphics," http://www.sun.com/960710/feature3/sketchpad.html. Sun Microsystems, 1994-1999.

150. R. Sternberg, "Intelligence and Lifelong Learning: What's New and How Can We Use It?" *American Psychologist*, vol. 52, no. 10, pp. 1134-1139, 1997.

151. R. Stoakley, J. Matthew, J. Conway, and R. Pausch, "Virtual Reality on a WIM: Interactive Worlds in Miniature," in *Human Factors in Computing Systems CHI 1995*. Denver, CO USA, May, 1995.

152. S. Sutherland, *The International Dictionary of Psychology*. New York: The Crossroads Publishing Company, 1996.

153. S. Sutphen, Private Communication. University of Alberta, 1999-2002.

154. S. Sutphen, E. Sharlin, B. A. Watson, and J. Frazer, "Reviving a Tangible Interface Affording 3D Spatial Interaction," http://www.cs.nwu.edu/~watsonb/school/docs/ski00.segal.pdf, in *Proc. 11th Western Canadian Computer Graphics Symposium*. Panorama, BC, Canada, March, 2000, pp. 155-166.

155. H. Suzuki and H. Kato, "Interaction-Level Support for Collaborative Learning: AlgoBlock - An Open Programming Language," in *Proceedings of CSCL*, 1995.

156. K. G. F. Thomas, H. E. Laurance, S. E. Luczak, and Jacobs W J., "Age-Related Changes in a Human Cognitive Mapping System: Data from a Computer-Generated Environment," *CyberPsychology and Behavior*, vol. 2, pp. 545-566, 1999.

157. P. W. Thorndyke and B. Hayes-Roth, "Differences in Spatial Knowledge Acquired from Maps and Navigation," *Cognitive Psychology*, vol. 14, pp. 560-589, 1982.

158. E. C. Tolman, "Cognitive Maps in Rats and Men," http://psychclassics.yorku.ca/Tolman/Maps/maps.htm, *The Psychological Review*, vol. 55, no. 4, July, pp. 189-208, 1948.

159. ToyLab, "Tonka Workshop," http://www.toylab.com/reviews/tonkaworkshop.html, 1999.

160. B. Ullmer, "Tangible Interfaces for Manipulating Aggregates of Digital Information," www.media.mit.edu/~ullmer/thesis. Ph.D. Dissertation. MIT, 2002.

161. B. Ullmer and H. Ishii, "The MetaDesk: Models and Prototypes for Tangible User Interfaces," in *Proceedings of UIST 1997*, October, 1997, pp. 1-10.

162. B. Ullmer and H. Ishii, "Emerging Frameworks for Tangible User Interfaces," in *Human Computer Interaction in the New Millennium*, J. M. Carroll, Ed. Addison-Wesley, 2001, pp. 579-601.

163. B. Ullmer, H. Ishii, and D. Glas, "MediaBlocks: Physical Containers, Transports, and Controls for Online Media," in *Proceedings of SIGGRAPH 98*. ACM, 1998, pp. 379-386.

164. M. Umaschi, "Soft Toys with Computer Hearts: Building Personal Storytelling Environments," in *Human Factors in Computing Systems CHI 1997; Electronic Publications: Demonstrations*. ACM, 1997, pp. 20-21.

165. J. Underkoffler and H. Ishii, "Illuminating Light: An Optical Design Tool with A Luminous-Tangible Interface," in *Human Factors in Computing Systems CHI 1998*, April, 1998, pp. 542-549.

166. J. Underkoffler and H. Ishii, "Urp: A Luminous-Tangible Workbench for Urban Planning and Design," in *Human Factors in Computing Systems CHI 1999*, May, 1999, pp. 386-393.

167. S. G. Vandenberg and A. R. Kuse, "Mental Rotations: A Group Test of Three-Dimensional Spatial Visualization," *Perceptual and Motor Skills*, vol. 47, pp. 599-604, 1978.

168. D. Waller, "Factors Affecting the Perception of Interobject Distances in Virtual Environments," *Presence: Teleoperators and Virtual Environments*, vol. 8, no. 6, December, pp. 657-671, 1999.

169. D. Waller, E. Hunt, and D. Knapp, "The Transfer of Spatial Knowledge in Virtual Environment Training," *Presence:*

Teleoperators and Virtual Environments, vol. 7, no. 2, April, pp. 129-143, 1998.

170. Y. Wang, A. Biderman, B. Piper, C. Ratti, and H. Ishii, "SandScape," http://tangible.media.mit.edu/projects/SandScape/SandScape.htm.
171. S. Ward, K. Abdalla, R. Dujari, M. Fetterman, F. Honoré, R. Jenez, P. Laffont, K. Mackenzie, C. Metcalf, M. Minsky, J. Nguyen, J. Pezaris, G. Pratt, and R. Tessier, "The NuMesh: A Modular, Scalable Communications Substrate," in *Proceedings of the International Conference on Supercomputing,* July, 1993.
172. S. Weiss, "Ivan Sutherland," http://www.digitalcentury.com/encyclo/update/sutherland.html. Jones Telecommunications & Multimedia Encyclopedia, 1994-1999.
173. P. Wellner, "The DigitalDesk Calculator: Tangible Manipulation on a Desk Top Display," in *Proceedings of UIST 1991.* ACM, 1991, pp. 27-34.
174. P. Wellner, "Interacting with Paper on the DigitalDesk," *Communications of the ACM,* vol. 36, no. 7, August, pp. 87-96, 1993.
175. M. Yim, "Digital Clay," http://www2.parc.com/spl/projects/modrobots/lattice/digitalclay. Xerox, PARC, January, 2002.
176. M. Yim, "A Self-Assembling Lattice Reconfiguration Robot," http://www2.parc.com/spl/projects/modrobots/lattice/telecube/index.html. Xerox, PARC, January, 2002.
177. M. Yim, "A Self-Assembling Lattice Reconfiguration Robot," http://www2.parc.com/spl/projects/modrobots/lattice/proteo/index.html. Xerox, PARC, January, 2002.
178. A. Zadorozny, Private Communication. University of Alberta, 2002.

Appendix A
Cognitive Cubes—Experimental Issues

Cognitive Cubes—Letter of Information

We are developing a device for assessment of early Alzheimer disease. Our device is called Cognitive Cubes. The device enables a user to build structures. It also enables a computer to sample or read these structures. People with early Alzheimer find it harder to construct structures. We hope that using our device a computer can automatically assess constructional abilities. If successful, Cognitive Cubes could be very useful. They might be used as an easy and handy assessment tool for Alzheimer disease. We believe Cognitive Cubes can hold great future benefit for caregivers in need of assessment tools for these diseases.

Today you can help us with one of the first experiments of our device. We are very grateful for your help! We will start by learning how to use Cognitive cubes. As you will see these cubes are easy to use. Later we will show you simple structure and ask you to build it, step by step, using Cognitive Cube. The structure will be displayed using a computer screen. You will have plenty of time to practice that. Finally, we will show you a structure and ask you to build it using Cognitive Cubes. This time you will be asked to build these structures as precisely and fast as you can. Cognitive Cubes will record all of your actions so we can evaluate them later. At the end of the assessment session we will interview you. We would like to learn more of what you thought of Cognitive Cubes.

Cognitive Cubes is a completely safe system. There are no risks or hazards in today's experiment. However, if you would like to stop the experiment at any point you are very welcome to do so. Please just indicate that you want to stop and we will do so immediately. You can decide not to perform any of the tasks or not to answer any of the tests questions. You can withdraw the experiment if you would like to at any point.

All your personal information will be kept strictly confidential. Your identity will not be published. We will analysis the results of your interaction tests without any link to your identity. We are hoping to publish the results of this work as an academic paper. If you would like us to send you a copy of such future paper please let us know and we would be happy to do so.

All information will be held confidential (or private), except when professional codes of ethics or legislation (or the law) requires reporting.

The information you provide will be kept for at least five years after the study is done. The information will be kept in a secure area (i.e. locked filing cabinet). Your name or any other identifying information will not be attached to the information you gave. Your name will also never be used in any presentations or publications of the study results.

The information gathered for this study may be looked at again in the future to help us answer other study questions. If so, the ethics board will first review the study to ensure the information is used ethically.

If you have any further concerns and you would like to contact someone who is not affiliated with our research team, please contact the University of Alberta Health Research Ethics Board (Karen Turpin, 492-0839).

Thank you very much again for volunteering to help us!

Ehud Sharlin and the Cognitive Cubes team members

Cognitive Cubes—Consent Form

Part 1: Researcher Information		
Name of Principal Investigator: Ehud Sharlin Affiliation: Department of Computing Science, University of Alberta Contact Information: 492-7418		
Name of Co-Investigator/Supervisor: Dr. Benjamin Watson Affiliation: Department of Computer Science, Northwestern University Contact Information: 1 847 491 3710		
Part 2: Consent of Subject		
	Yes	**No**
Do you understand that you have been asked to be in a research study?		
Have you read and received a copy of the attached information letter?		
Do you understand the benefits and risks involved in taking part in this research study?		
Have you had an opportunity to ask questions and discuss the study?		
Do you understand that you are free to refuse to participate or withdraw from the study at any time? You do not have to give a reason and it will not affect your care.		
Has the issue of confidentiality been explained to you? Do you understand who will have access to your records/information?		
Do you want the investigator(s) to inform your family doctor that you are participating in this research study? If so, please provide your doctor's name: ______________________________ (This question is optional).		
Part 3: Signatures		
This study was explained to me by: ______________________________ Date: ______________________________		
I agree to take part in this study. Signature of Research Participant: ______________________________ Printed Name: ______________________________		
Witness (if available): ______________________________ Printed Name: ______________________________		
I believe that the person signing this form understands what is involved in the study and voluntarily agrees to participate. Researcher: ______________________________ Printed Name: ______________________________		
* A copy of this consent form must be given to the subject.		

Cognitive Cubes—Protocol

<Remarks in brackets are directed for the administrator only>

1. **Today is:** **The experiment takes place in:**
2. **Verify constant physical conditions:**

Experiment Constant Physical Conditions	
Screen diagonal: 125cm	Distance to screen: 186cm
Interaction (table) height: 84 cm	Bright light conditions
Subject is seated in front of the screen. All the blocks, including base cube, are presented to the subject.	

3. **Introduction:**

"Hello, today we will perform an experiment of a new system called Cognitive Cubes. Cognitive Cubes are a set of computerized cubes that can be connected to a computer. When you build with cognitive cubes our system 'knows' which structure you have built. We want to know if the general public can use Cognitive Cube for all kind of different tasks including cognitive assessment tasks. This is why we asked for your help. Before we start our cognitive cubes experiment I will read out loud our information letter and ask you to sign the consent form."

4. **Read out loud information letter.**
5. **Ask subject to read and to sign the consent form.**
6. **Personal information**

"I would like to write down a few details about you."

"First I will write down your name, your age and your contact address/email."

"Since we would never use your real name with the results, due to confidentiality, we would like to choose a nickname for your experiment. You can choose a nickname or let us choose one for you."

"Now I would like to ask you a few questions about your background and education"

"Let us start by your occupation."

"Now please let me know how many years of formal education you had and towards which degrees/certificates you studied."

"I would like you to tell me if you ever, even as a child, played with puzzles, constructions sets, Lincoln Logs, Legos, technical Legos, etc. "

"Do you have any experience in 3D design, technical construction, sculpturing, etc.?"

"Also please let me know if you ever played video or computer games."

"Last, please let me know, without getting into details, what is your health condition"

<u>Subject's Info</u>

- Subject's Name: Age:
- Subject's contact address/Email:
- Subject's system confidential nickname:
- Subject's education:

  Years of education: High school/certificates/degrees:
- Puzzles, Construction sets, Lincoln Logs, Lego, technical experience):
- Experience in 3D design, technical construction, sculpturing, etc.:
- Experience in video games and/or computer games:
- Health condition:

7. **<u>Administer mental status and right handed test:</u>**

"We will start the experiment with two general tests. I am going to ask you a few questions, please try to answer them as best as you can."

Administer mini mental test

Administer right-handed test

8. Welcome Session

"Please remember that if you feel tired or want to stop the experiment for any reason at any point you just need to say so, and we will stop or finish the experiment".

"Now we can start working with Cognitive Cubes. "

"In all the Cognitive Cubes tasks we will ask you to try and build, using the cubes, a shape that will be shown on the screen, while trying to be as precise as possible."

"We will start by learning how to connect and disconnect Cognitive Cubes in a session that is called Welcome. This session results will not be used for assessment."

"Let us start with welcome1. You should try and match the shape that is presented on the screen precisely, or reach the best match you can. While you build with the cubes our system records your actions and the time it took you to complete the task."

<During Welcome1 and Welcome2 the administrator should hold a set of physical cubes and demonstrate the instructions using them>

"When you want to connect two Cognitive Cubes you should connect any two faces of the two cubes. You can do it by connecting the male connectors on one face to the female connector on the other face <Show to Subject>. You can also match the two blue stripes on each face and then connect the faces. If one stripe 'kisses' the other the cubes will connect."

"In order for the computer to "know" that you made the connection, all the cubes should be connected to the "base cube", the cube that has a wire connected to it. Let us try to perform this".

<Supervise connection to base cube, wait for audio feedback>

"You can hear the connection sound. You will hear this sound whenever you will connect a new cube to your structure"

"Let us now try and disconnect the cubes"

"You can hear the disconnection sound. You will hear this sound whenever you will disconnect a cube from your structure"

"Let us now practice connection and disconnection of cubes".

<Practice in front of Welcome1, the subject should perform several connections and disconnections of the two cubes. >

<Possibly, only if the need arise, mention the following issues: hands on blanket: “Please try and keep the cubes on top of the blanket, this will help you build complex structures”; Making a good mechanical connection: “The cubes will connect better to each other if you will snap them firmly together”. Delayed sound feedback: ”Please do not worry if you do not hear the sound right after a connection. Please just keep on working. Our system has short delays sometimes.>.

“In order to finish the Welcome1 task, you should now connect a cube to base cube. “

“As long as the cubes shape match the shape on the screen, you shouldn’t worry about the direction and orientation of the structure” <Explain how the orientation changes without affecting the shape by rotating the physical structure>

“As I mentioned before, you should try to be precise and accurate when you try to match the physical structure. We will not tell you if you are correct or not – it is up to you try to do the best you can, match the structure on the screen and tell us if you think you are done. Our system records all the moves you are making and the time it took you to make them.”

“Let us know when you think you finished Welcome1”

“Good job!”

“Are you ready for the next task?”

“Let us try Welcome2. Remember, try to be as accurate as possible and remember our system is recording your time”.

“As you see, the system shows you another cube after you connected the first one”

“Let us know when you think you finished Welcome2”

“Good job!”

“Are you ready for the next task?”

“Let us try Welcome3. Remember, try to be as accurate as possible and remember our system is recording your time”.

“Again as you see, the system shows you another cube after you connected the first one”

“Let us know when you think you finished Welcome3”

“Good job!”

“We ask you to always connect only a single cube at a time to the structure. However, you can disconnect multiple cubes at once from your structure if you need to.”

<Follow the same routine for the other Welcome tasks>

<In case of difficulty use: "Take your time", "Do the best you can">

9. Follow Me Session

"Like the Welcome tasks, the Follow Me tasks will be used for training. The tasks will present you the structure step by step, asking you to build it using the cubes. The tasks might be a bit more difficult than the Welcome tasks."

"Remember, we ask you to always connect only a single cube at a time to the structure. However, you can disconnect multiple cubes at once from your structure if you need to."

"Are you ready for Follow Me1?"

"Let us start Follow Me1. Remember, try to be as accurate as possible and remember our system is recording your time".

"Let us know when you think you finished Follow Me 1"

"Good job!"

"Are you ready for the next task?"

<Follow the same routine for the other Follow Me tasks>

<5 MIN. BREAK>

10. Match Me Session

"The Match Me tasks will be used for assessment. The system will present you the complete structure at once (!), asking you to build it using the cubes."

"As I mentioned before, you should try to be precise and accurate when you try to match the physical structure. We will not tell you if you are correct or not – it is up to you try to do the best you can, match the structure on the screen and tell us if you think you are done. Our system records all the moves you are making and the time it took you to make them."

"Remember, we ask you to always connect only a single cube at a time to the structure. However, you can disconnect multiple cubes at once from your structure if you need to."

"Are you ready for Match Me 1?"

"Let us start Match Me1. Remember, try to be as accurate as possible and remember our system is recording your time".

"Let us know when you think you finished Match Me 1"

"Good job!"

"Are you ready for the next task?"

<Follow the same routine for the other Match Me tasks>

<Possibly, only if the need arise, mention the following issues: Delayed sound feedback: "Please do not worry if you do not hear the sound right after a connection. Please just keep on working. Our system has short delays sometimes.">.

<In case of difficulty use: "Take your time", "Do the best you can">

11. Reshape Me Session

"The Reshape Me tasks will be used for assessment. Again the system will present you the complete structure at once, asking you to build it using the cubes. However, this time you will start from a ready structure built by me from the cubes. You will need to reshape this existing structure by removing and adding cubes to it, trying to match the structure presented on the screen. You can always add and remove cubes from the system, but you should try and make as few step as possible."

"Remember, we ask you to always connect only a single cube at a time to the structure. However, you can disconnect multiple cubes at once from your structure if you need to."

"Are you ready for Reshape Me 1?"

"Let us start Reshape Me1. Remember, try to be as accurate as possible and to use as few steps as possible. Our system is recording your time and counting your steps".

"Let us know when you think you finished Reshape Me 1"

"Good job!"

"Are you ready for the next task?"

<Follow the same routine for the other Reshape Me tasks>

<Possibly, only if the need arise, mention the following issues: Delayed sound feedback: "Please do not worry if you do not hear the sound right after a connection. Please just keep on working. Our system has short delays sometimes.">.

<In case of difficulty use: "Take your time", "Do the best you can">

Semi Structured Interview

What was your general experience (fun? Hard?):

On a scale of 1 to 10, how difficult was the whole experiment?

On a scale of 1 to 10, how difficult was the most difficult task?

How did you find the tasks? Welcome? Follow-Me? Match-Me? Reshape Me? Which tasks were the hardest?

Can you give me feedback on the cubes? (size , weight, color)

How did you connect the cubes? Did you use the clips or the color stripes?

Can you comment on the display?

Can you comment on the rotation?

Can you comment on the sound?

If we were to work with more mature adults what would you suggest we do in order to improve our system and tasks?

Admin Comments

Appendix B

Cognitive Map Probe—Experimental Issues

Cognitive Map Probe—Letter of Information

We are interested in knowing whether computers can help us evaluate how people find their way around new places. The computer program we have developed is called the "Cognitive Map Probe" or "CMP". Using this program, we can represent a "model" or small version of a neighborhood. When we test how a person finds his or her way around the model neighborhood, the computer can record the person's progress. If this way of evaluation is acceptable to people, then we may be able to use it for evaluating persons with memory problems that make it difficult for them to find their way around new places.

We will start the experiment by asking you a few questions about your education. We will also ask about your experience with computers and with computer games. We might also ask you to perform a short memory test. You will then learn how to use the CMP. We will show you a simple neighborhood and ask you to build it, step by step. The neighborhood you build will be shown on a large screen. When you are comfortable using the CMP, we will present you with a more complex neighborhood and ask you to build it on the CMP. We will repeat this with several neighborhoods. The CMP will record your progress. At the end of the session we will ask you what you liked and disliked about the CMP. The entire experiment should take about one hour. You are welcomed to take a brake whenever you want to.

The CMP is a completely safe system. There are no risks or hazards in today's experiment. However, if you would like to stop the experiment at any time, you may do so. You can decide not to perform any of the tasks or not to answer any of the questions. You can withdraw from the experiment at any point if you want to.

All your personal information will be kept strictly confidential. Your identity will not be published. We will analyze the results of your performance without any link to your identity. We are planning to publish the results of this work as an academic paper. If you would like us to send you a copy of such a future paper please let us know and we will do so.

The information you provide will be kept for at least 5 years after the study is done. The information will be kept in a secure locked filing cabinet. Your name or any other identifying information will not be attached to the information you gave. Your name will also never be used in any presentations or publications of the study results. All information will be kept private expect when codes of ethics or the law requires reporting.

The information gathered for this study may be looked at again in the future to help us answer other study questions. If so, the ethics board will first review the study to ensure the information is used ethically.

Thank you very much,

Ehud Sharlin and the Cognitive Map Probe team members

Cognitive Map Probe—Consent Form

Part 1: Researcher Information		
Name of Principal Investigator: Ehud Sharlin Affiliation: Department of Computing Science, University of Alberta Contact Information: 492-7418		
Name of Co-Investigator/Supervisor: Dr. Benjamin Watson Affiliation: Department of Computer Science, Northwestern University Contact Information: 1 847 491 3710		
Part 2: Consent of Subject		
	Yes	**No**
Do you understand that you have been asked to be in a research study?		
Have you read and received a copy of the attached information letter?		
Do you understand the benefits and risks involved in taking part in this research study?		
Have you had an opportunity to ask questions and discuss the study?		
Do you understand that you are free to refuse to participate or withdraw from the study at any time? You do not have to give a reason and it will not affect your care.		
Has the issue of confidentiality been explained to you? Do you understand who will have access to your records/information?		
Part 3: Signatures		
This study was explained to me by: __________ Date: __________		
I agree to take part in this study. Signature of Research Participant: __________ Printed Name: __________		
Witness (if available): __________ Printed Name: __________		
I believe that the person signing this form understands what is involved in the study and voluntarily agrees to participate. Researcher: __________ Printed Name: __________		
* A copy of this consent form must be given to the subject.		

Cognitive Map Probe—Protocol

<Remarks in brackets are directed to the administrator only>

1. **Today is:** **The experiment takes place in:**
2. **Verify static physical conditions:**

Experiment Constant Physical Conditions

Screen diagonal: 205cm Distance projector-screen: 305cm

Screen vertical (top and bottom, measured from floor): 125-250cm

CMP (table) height: 72 cm

Mildly-dim light conditions (main light- off, table lamp and floor lamp – on).

Subject is seated in front of screen, moving chair next to the CMP, not allowed to walk around the CMP interface.

CMP room 3-57 layout:

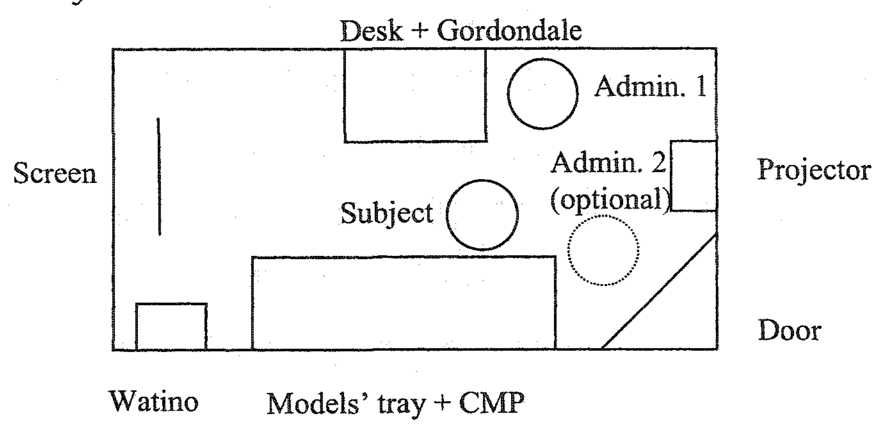

<Detailed procedure – revision for Admin.>

1. Tangible Start/End: (a) When subject is ready to start, insert and remove the start model (b) When subject finishes the interaction, insert the start model, and leave it in. (c) Disconnect all the models from the CMP (d) Remove the start model (e) For another scene go to 'a'. (f) To terminate the experiment connect/disconnect start model 3 times.

2. Remember to position the physical compass and the bus in all training scenes. In assessment scenes, remove the bus and leave the compass.
3. When the experiment has finished – verify the backup of CMP & world results file.
4. Document every event of unrecognized object (UFO), or mismatch error, using the UFO Number, name and orientation (use last page).
5. Verify server-client working fine.

3. <u>Introduction:</u>

"This is an experiment that uses special models of houses and buildings connected to a computer to measure how well people learn their way in new places. <point to models> When you connect the house models to a special board the computer can record your actions <point to board>. We want to know how useful these models are for assessing people who have memory problems. <say this to the healthy subjects> We also want to know how acceptable the models are to the general population. This is why we asked for your help. Before we start the experiment I will read out loud our Information Letter and ask you to sign the Consent Form."

4. <u>Read out loud Information Letter.</u>

5. <u>Ask subject to read and to sign the Consent Form.</u>

6. <u>Personal information</u>

"I will write down your name, your age and your contact address/email."

<Assign nickname= date + counter + experimental flavor (CMP/WIMP/PAPERPEN) +, e.g., third subject for the day, performing PAPERPEN on July 7th is: JULY7_3_PAPERPEN>.

<FOR AD only: perform MiniMental, Assess level of AD, if using Aricept ask for duration of use>

Subject's Info

•	Date
•	Location of experiment: CMP Room (Other ___________)
•	Subject's Name: Age:
•	Administrator's name
•	mailing address/Email:

•	Subject's system confidential nickname:
•	<AD only>: Minimental: AD level: Aricept use [Months]:

7. **Training I - Physical**

<Setup: Virtual OFF; Physical setup includes: Scenario #1 ON, Compass ON, Bus ON, CAMERA ROLL, Say out loud ONE>

"If you feel tired, tell me and we will take a break. You may also stop the experiment at anytime."

"Let's start by training for our exercise"

<Point to physical setup>

"This is a typical neighborhood in Edmonton in small scale. The compass shows you where North is, it will point to the same direction throughout our experiment. The bus shows you where you are starting your tour of the neighborhood. Let's imagine that you are a passenger in this bus. I <or Admin #2, if applicable> will be the bus driver. The route that we are about to take will be repeated in all our bus rides today. The neighborhoods will change, each containing different houses, but the bus route will stay the same"

"The key to this exercise is to try to remember the location and orientation of the landmarks you are visiting in the neighborhood."

"Let us take the bus ride and learn the route."

<Physically, slowly move bus on interaction route>

"As we drive forward we can see the orange church to our right, here we can see the tower (steeple) and the cross. We now turn right at this intersection and the church is on our right. We drive forward to the end of the street. We make a U turn and drive forward again. We then turn right into this side street. At the end we make another U turn, and turn right at the intersection. The orange church is now on our left. In front of us to the right is the green residential house, here we can see the entrance to the house. We take another right and then U turn at the end. We keep on following the route, with the green house now on right. We make another U turn and then turn right. We end our journey at the same place we started it."

"The bus will take you on this same route throughout the experiment today"

Training II – Virtual

<Setup: Virtual ON Scenario #1 ON; Physical setup includes: Scenario #1 ON, Compass ON, Bus ON, Say out loud ONE >

"Let's continue training for our exercise"

"Let's visit the same neighborhood in Edmonton, now presented on the big screen"

<Point to screen>

"Note the sign welcoming you. We will see this sign whenever we start our drive".

" The compass shows you where North is. It shows the same direction as the compass on the table is showing"

"Let's imagine that you are a passenger in the bus we saw before. I <or Admin #2, if applicable> will be the bus driver. We start our journey exactly at the bus's location on the table".

"The key to this exercise is to try to remember the location and orientation of the buildings in the neighborhood you are visiting on the screen so you can reconstruct this neighborhood on the table. The computer records your decisions and the time it took you to make them. Try to be as accurate as you can. Try also to work as fast as you can.".

"Let us start our bus ride"

<Start Virtual Drivethrough>

"At any point you can ask the driver to stop the bus"

<Stop the Bus. Make sure that the church is visible>

<Test for basic transfer>

"Is the church you see on the screen and the church on the table the same? Are they located at the same location?"

<Wait for reply, if no reply, say "yes, they are the same. Can you see that?". Mention Tower /Steeple for recognition>

"At any point you can ask to continue driving"

<Continue driving>

"Or to look around by rotating in place"

<From driving, transfer to rotation directly>

<If subject is passive ask: "please let me know if you would like us to continue our drive">

<Continue driving>

"We drive forward to the end of the street. We make a U turn and drive forward again. We then turn right into this side street..

"At the end we make another U turn, and turn right at the intersection. The Orange Church is now on our left."

"In front of us to the right is the green residential house."

"Again, remember that at any point you can ask the driver to stop the bus"

<Stop the Bus. Make sure green house is visible>

<Test for basic transfer>

"Do you think the green house we saw on the screen is at the same place and orientation as the green house on the table?"

<Wait for reply, if no reply, say "yes, they are the same. Can you see that?">

"At any point you can ask to continue driving"

<Continue driving>

"Or to look around by rotating in place"

<From driving, transfer to rotation directly>

<If subject is passive ask: "please let me know if you will like us to continue our drive">

<Continue driving>

We take another right and then U turn at the end. We keep on following the route, the green house is on right. We make another U turn and turn right. We end our journey at the same place we started it."

"Note the sign saying that we have reached the end of our drive"

"The bus will take you on this same route throughout the experiment today, but each time different neighborhoods would be displayed"

<In case subject failed transfer questions, repeat drive, if unsuccessful after three repetitions - stop experiment>

Training III – Transfer Virtual to Physical

<Setup: Virtual ON Scenario #2 ON; Physical setup includes: Scenario #1 ON, Compass ON, Bus ON, Say out loud TWO >

"Let's continue our training"

"Your task is to try to remember the location and orientation of the buildings in the neighborhood you are visiting on the screen so you can reconstruct this neighborhood on the table. The computer records your decisions and the time it takes you to make them. Try to be as accurate as possible".

"Let's visit our neighborhood again"

"Tell me when you would like to start the drive, please let me know if there are any differences this time between the neighborhood on the screen and the one on the table".

<On participant request, start driving>

<If participant is passive remind her that: "At Any point you can ask the driver to stop the bus" "At any point you can ask to continue Driving" "Or to look around by rotating in place">

<If subject is passive, ask:>

"Are these buildings in the same place on the screen and on the table?"

<If subject is still passive, ask:>

"Can you see that on the screen the church moved from to the other side of the street? Can you show me where did it move to on the table?"

<When drivethrough is done, ask:>

"What do you think? Is the neighborhood you saw on the screen the same as the one on the table" "Which buildings were in different locations? Please show me on the table."

<Subject should point to the church's new position on the table, and be aware of new orientation >

<If subject's seems to understand the requirements, but faces problems locating the new position, refer to the virtual world again (still ON) or even repeat drivethrough>

<If participant was completely passive, and couldn't follow the transfer, stop the experiment >

"Good job!"

Training IV – Interaction

<Setup: Virtual OFF; Physical setup includes: Scenario #1 ON, Compass ON, Bus ON>

"Let's see how we can move the church into its new place"

"All you have to do is take the church out of its current location on the table"

<Take church out>

"And relocate it to its correct new location and orientation. When you insert the church into its new position the borders of the church's base should be aligned to the grid on the table <point to grid line>. You should then press the church into its new place, like this"

<Demonstrate connection of church to the position pointed by the participant earlier. Then place it back in original, erroneous, position>

"Now you do it, please."

"Good job!"

<If participant faces problems, repeat explanations. If persists, assist in final phases of guiding the connector into the slot>

Training V – Summary

<Setup: Virtual ON Scenario #3 ON; Physical setup includes: No Scenario (all models OFF), Compass ON, Bus OFF, Say out loud THREE>

"Let's try the last part of our training. This will follow the exact same routine of the exercise we will perform soon and will summarize what we just learned"

"Your task is to try to remember the location and orientation of the buildings in the neighborhood you are visiting on the screen so you can reconstruct this neighborhood on the table. The computer records your decisions and the time it took you to make them. Try to be as accurate and as you can".

"Let's drive through the neighborhood. Remember you can ask the bus driver to stop or rotate at any time".

<Perform drivethrough scenario #3 following subject's order>

<If facing drivethrough problems, repeat relevant parts of earlier training. If persists, stop experiment>

"Now let us try to rebuild the neighborhood that we just visited, on the table."

<Setup: Virtual OFF; Physical setup includes: No Scenario (all models OFF), Compass ON, Bus OFF>

"Try to place the correct building that you saw on the screen, in its correct location and orientation on the table"

<If facing transfer or interaction problems, repeat relevant parts of earlier training. If persists, stop experiment>

"Good job!"

Assessment I-X (repeat)

<Take physical models off table>

"We have only seven tasks like this ahead of us"

"If you feel tired, tell me and we will take a break. You may also stop the experiment at anytime."

<Setup: Virtual ON Scenario 4-13 ON; Physical setup includes: No Scenario (all models OFF), Compass ON, Bus OFF, Say out loud SCENARIO NUMBER >

"Your task is to try to remember the location and orientation of the buildings in the neighborhood you are visiting on the screen so you can reconstruct this neighborhood on the table. The computer records your actions and the time it took you to make them. Try to be as accurate as you can".

"Let's drive through the neighborhood. Remember you can ask the bus driver to stop or rotate at any time"

<Perform drivethrough of current scenario following subject's driving instructions"

<If facing drivethrough problems, repeat relevant parts of training. If persists, stop experiment>

"Now let us try to rebuild the neighborhood that we just visited, on the table."

<Setup: Virtual OFF; Physical setup includes: No Scenario (all models OFF), Compass ON, Bus OFF>

"Try and place the buildings that you saw on the screen, in their correct location and orientation on the table."

<If subject's facing difficulties and frustration, stop experiment. >

"Good job!"

Semi Structured Interview

Describe your general experience (fun? Hard? <Use only these words to lead if subject needs prompting>):

On a scale of 1 to 10, how difficult was the whole experiment <Use VAS line on next page>?

Did you use any special technique to remember/memorize the map?

Admin Comments

UFO events:

How Difficult was the experiment?

Mark an "X" on the line which indicates a range of feelings.

Not difficult at all | Extremely difficult

Appendix C
List of Abbreviations

HCI—Human Computer Interaction

TUI—Tangible User Interface

CMP—Cognitive Map Probe

AD—Alzheimer's Disease

DOF—Degrees Of Freedom

UI—User Interface

Standard UI—Standard User Interface

2D—Two-Dimensional

3D—Three-Dimensional

WIMP—Windows-Icon-Menu-Pointer

I/O—Input/Output

MRT—Mental Rotation Test

VR—Virtual Reality

VRSR—Virtual Reality Spatial Rotation